Preliminary Test Results for the SVX4

L. Christofek, K. Hanagaki, P. Rapidis, M. Utes

Abstract: We present and summarize the preliminary test results for SVX4 chip testing.

1 Introduction

There are presently two versions of the SVX4. Version 2 has on-chip bypassing and Version 1 does not. The on-chip bypassing is a layer of transistors under the front-end analog pipeline that acts as a bypassing capacitor for the voltage supply. Its size is about a microfarad. We aggressively choose to test Version 2 because of this feature.
Also, the SVX4 was designed to operate in two modes: D0 and CDF. One can set which mode the chip will operate by placing a jumper in the proper position on the SVX4 chip carrier.
In either mode, the chip can either use the operating parameters from the shift register or the shadow register. Similarly, this is selected by placing a jumper on the SVX4 chip carrier. This chip has this feature because it was unknown whether the new design of the shadow register would be operable. The shadow register is also call the SEU register or Single Event Upset register.

2 Experimental Setup

Most of the results reported in this paper come from the Stimulus Test Stand. The test stand is presently located at Wilson Hall on the fourteenth floor. A detailed report on the Stimulus Test Stand and how it works in given in [1].

3 Gain

In order to measure the gain, we wire-bonded four channels of the chip to four different valued external capacitors. We then injected an external pulse into these channels using a Hewlett-Packard 8112A Pulse Generator. The external pulse passes through a variable attenuator to reduce its amplitude into the μV range (larger pulses will destroy the amplifier input and channel). We can control the pulse height of the signal in two independent methods: 1) varying the pulse height of the pulse generator (mV~V) and 2) varying the attenuator (10~100 dB). The injected charge is then calculated by using

$$Q = C_{ext} \times V_{pulser} \times 20 \log_{10} (V_{cap}/ V_{pulser})$$

where Q is the total charge injected, C_{ext} is the external capacitance bonded to the channel, and $20 \log_{10} (V_{cap}/V_{pulser})$ is the attenuation factor [dB] with V_{cap} being the actual voltage the external capacitor sees and V_{pulser} being the output voltage of the pulser. After injecting the pulse into the channel, we measure the ADC output of the channel. We then take the difference, Δ, between the ADC output of the injected signal and the ADC output of the pedestal. The gain, G, is then calculated by

$$G \equiv Q / \Delta .$$

In Figure 1, we give a cartoon schematic to show the experimental setup.

[1] D0Note xxx, The Stimulus Test Stand, L. Christofek, P. Rapidis, A. Reinhard

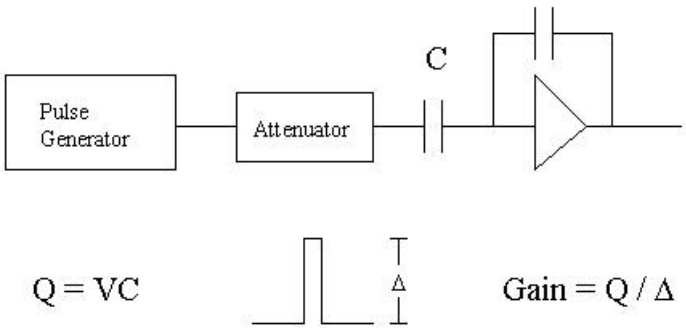

Figure 1 Cartoon schematic of the experimental setup used to measure the gain of the SVX4 chip. An external pulser generates a pulse that passes through an attenuator, which is then injected into an external capacitor wire-bonded to a single channel of the SVX4. The charge injected can be calculated from Q=VC. The gain is then calculated by measuring the difference between the charge injected signal and pedestal and dividing that quantity into the total charge injected. This gives the number of electrons per ADC count.

Because the SVX4 uses double correlated sampling, it is very important to setup the correct timing of the injected pulse. In simple terms, the preamplifier output is sampled during the falling edge of the front-end clock and then sampled again at the next rising edge of the front-end clock and the difference is taken as the signal to be digitized (this is the operational definition of double correlated sampling). It is therefore important to inject the external pulse directly after the falling edge of the front-end clock to allow the maximum integration time possible.[2] We show the results in Table 1 and Table 2.

BW\C	10 pF	33 pF	68 pF
0	706	899	1332
6	733	1119	1756
15	860	1561	2575

Table 1 Measured gain in terms of the number of electrons per ADC count. We measured the gain as a function of three bandwidths. The above data is for a Version 2 chip is d0 mode. This data was taken using a ramp slope setting equal to 1.

BW\C	10 pF	33 pF	68 pF
0	682	1023	1246
6	706	1262	1756
15	824	1854	2146

Table 2 Measured gain in terms of the number of electrons per ADC count. We measured the gain as a function of three bandwidths. The above data is for a Version 2 chip in CDF mode. This data was taken using a ramp slope setting equal to 1.

[2] This time is dependent on the front end clock frequency and duty cycle. Two other factors must be taken into account when calculating the gain: 1) the frequency of digitization and 2) the slope of the ramp. Increasing the digitization frequency decreases the gain. Similarly, increasing the slope of the ramp decreases the gain.

The gain can also be measured by using the internal calibration injection circuitry. By injecting a calibration pulse on a channel that is not bonded out through an external capacitor, the number of ADC counts per mV can be calculated for a specific bandwidth setting. Assuming an input capacitance of 25 fF for the calibration injection capacitor, we calculated for the gain a value of 591 electrons per ADC count. The results are shown in Table 3. This is about a 13% difference from using the external pulser method.

mV	0	30	77	125
ADC counts	142.8	150.8	163.5	175.2
Δ	0	8.0	20.7	32.4
mV/Δ	-	3.75	3.72	3.86

Table 3 Measured values for the ADC output as the calibration injection voltage is varied. The value of ramp slope was set to 1 and the bandwidth was set equal to 0. The difference between the any two values can be used to calculate the value of mV per ADC count. The average value is 3.78. The above data is for a Version 2 chip in CDF mode.

4 Noise

Using the measured gain from above, the RMS of the pedestal can be converted into the number of electrons also known as the equivalent noise charge (ENC). We show this in Figure 2. This measurement includes contributions from the front and back-end of the chip as well as any common mode noise contributions that may be in the system. We then fitted the data using the data analysis package in EXCEL. The results are shown in Table 4.

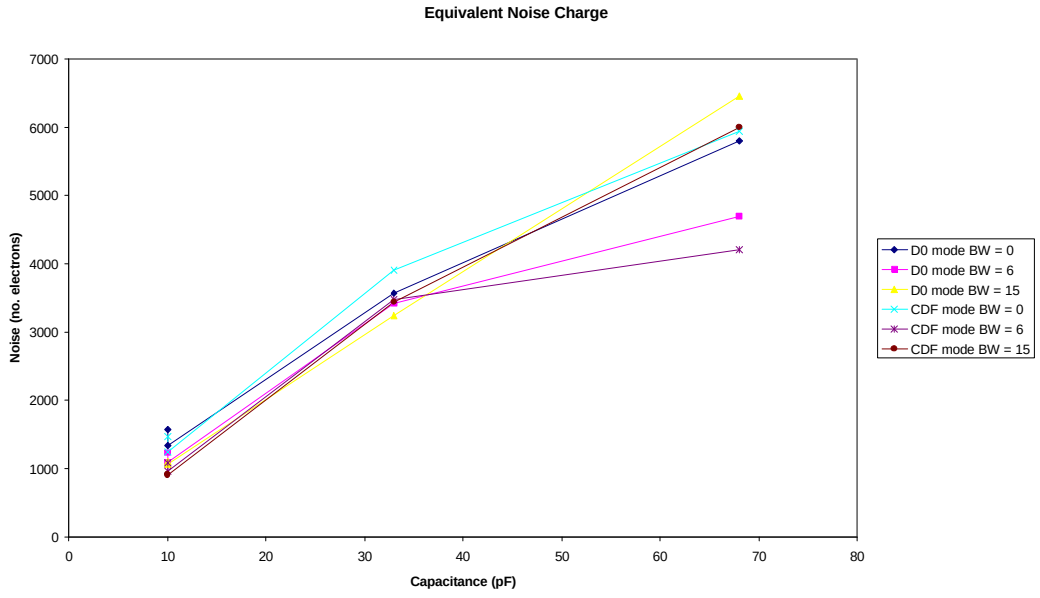

Figure 2 Plot of the equivalent noise charge as a function of capacitance in both D0 mode and CDF mode for three different bandwidth settings. We used a Version 2 chip with three different valued capacitors (10 pF, 33 pF, and 68 pF).

	D0 mode	CDF mode
BW = 0	715.9 + 80.08 × C [pF]	786.8 + 75.51 × C [pF]
BW = 6	728.5 + 56.41 × C [pF]	736.9 + 62.00 × C [pF]
BW = 15	140.9 + 88.48 × C [pF]	166.8 + 92.65 × C [pF]

Table 4 The ENC for bandwidths 0, 6, and 15 obtain from fitting a straight line to the data in Figure 4 in both D0 mode and CDF mode.

5 Rise Time

This current chip submission has several test pads on the chip for detailed studies. Channel 127 has a special probing pad on top of the chip in which we can study the preamp output. This pad allows the measurement of the rise time of the preamplifier to be made. We used a P6243 Tektronix FET probe. It is important to note that the output buffer has to be correctly biased before the output can be measured. We biased the pmos output using a 5.92V power supply through a 1 kΩ resistor.
We used a Version 2 chip in CDF mode to measure the rise times using an external calibration injection pulse with a magnitude of 2.4 V and a large pulse width (>250 ns). The injection capacitance for the channel was 10 pF. The output from the preamplifier is shown in Figure 3. The results for the rise time using the 10-90% definition are given in Table 5.

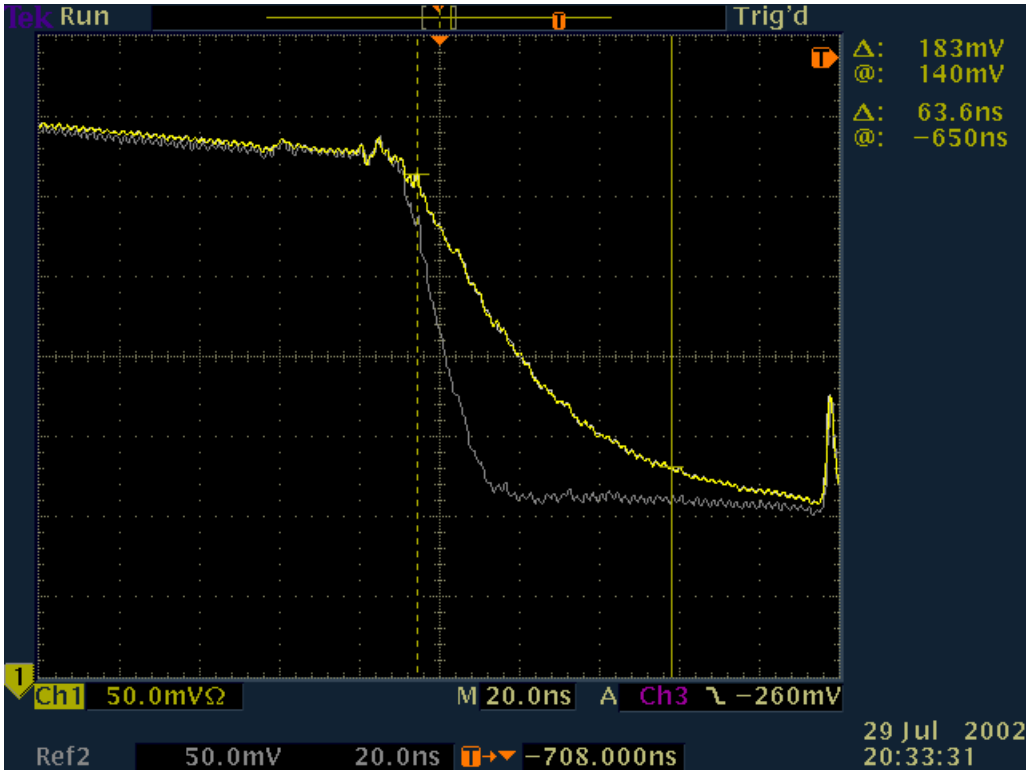

Figure 3 A picture of the preamplifier output showing the rise time of channel 127. The yellow line corresponds to a bandwidth setting of 15 and the white line corresponds to a bandwidth of 0. The spike at the far right of the picture is pick-up from the front-end clock.

Bandwidth	10 pF load
0	20 ns
3	32 ns
7	46 ns
11	58 ns
15	64 ns

Table 5 The measured rise times of the preamplifier using the 10-90% definition of rise time for various bandwidth settings. They were made using channel 127.

We also used this same configuration as above to measure the reset time of the preamplifier. In Figure 4, we see the reset of the preamplifier. The reset has a slight over shoot, but the value we measure for the reset time is 17.97 ns.

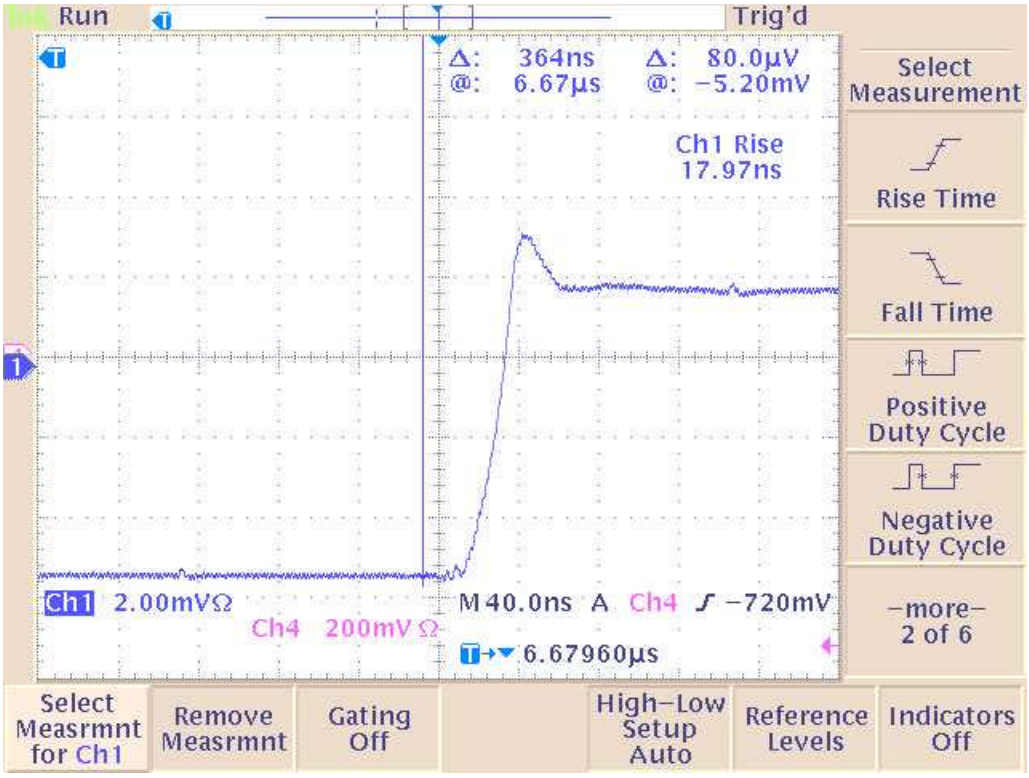

Figure 4 Screen shot from an oscilloscope showing the preamplifier reset in action. This was taken from channel 127.

6 Probing the Ramp

The ramp can also be investigated because it has test pads that allow direct observation. In Figure 5, we show the actual ramp in action from a Version 2 chip. We measured the ramp slope using the FET probe mentioned above. Table 6 shows the measured slope and the design values. In the design, the slope is determined from

$$0.5\ \text{mV/ns} \times [\ 1 + (\ 2 \times R0\ ) + (\ 2 \times R1) + (\ 1 \times R2\ )\ ]^{-1}$$

where R0, R1, and R2 correspond to each bit setting.
We confirmed the functionality of the ramp pedestal settings (a four bit field) by downloading all values of the pedestal and then measuring the pedestals using the chip ADC. We did the analogous thing for the ramp range settings using a single ramp pedestal setting. We show the output of the ADC for both studies in Figure 6.

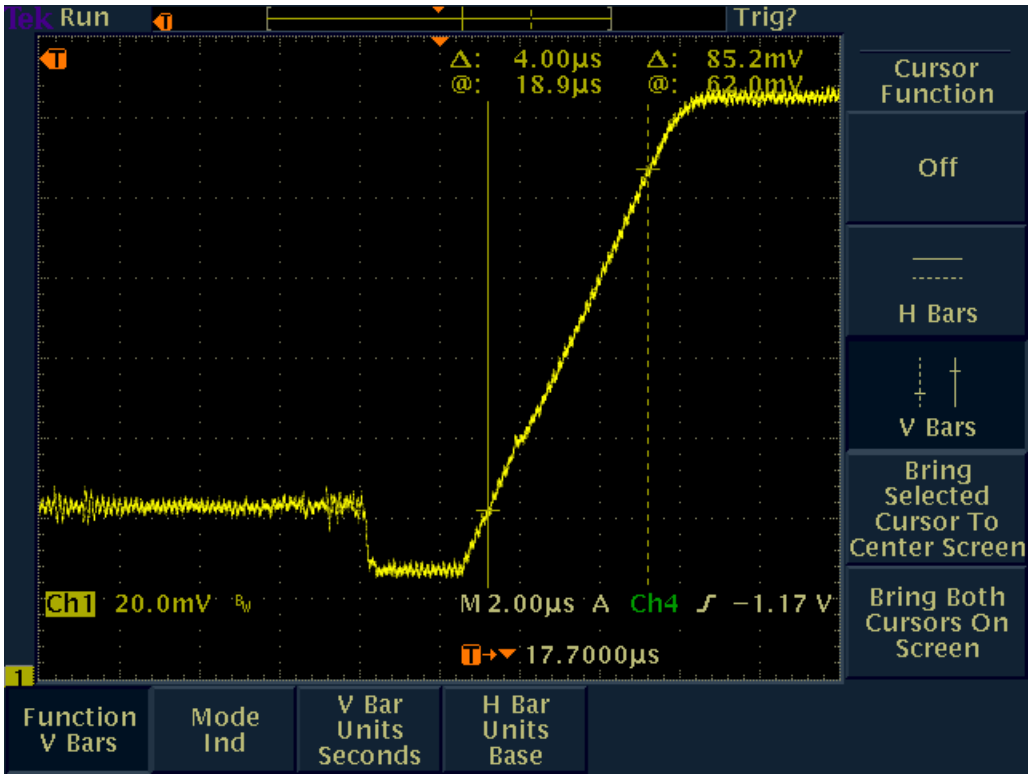

Figure 5 A picture showing the actual ramp for a Version 2 chip. We used the oscilloscope's functions to measure the ramp slope.

Ramp Range Setting	Bit Pattern	Design (mV/ns)	Measured (mV/ns)
0	000	0.500	0.416
1	001	0.167	0.149
2	010	0.167	0.137
3	011	0.083	0.082
4	100	0.250	0.337
5	101	0.125	0.111
6	110	0.125	0.105
7	111	0.083	0.073

Table 6 Ramp Range settings with the bit pattern - listed in R2, R1, R0 order - along with the design and the measured values for the ramp.

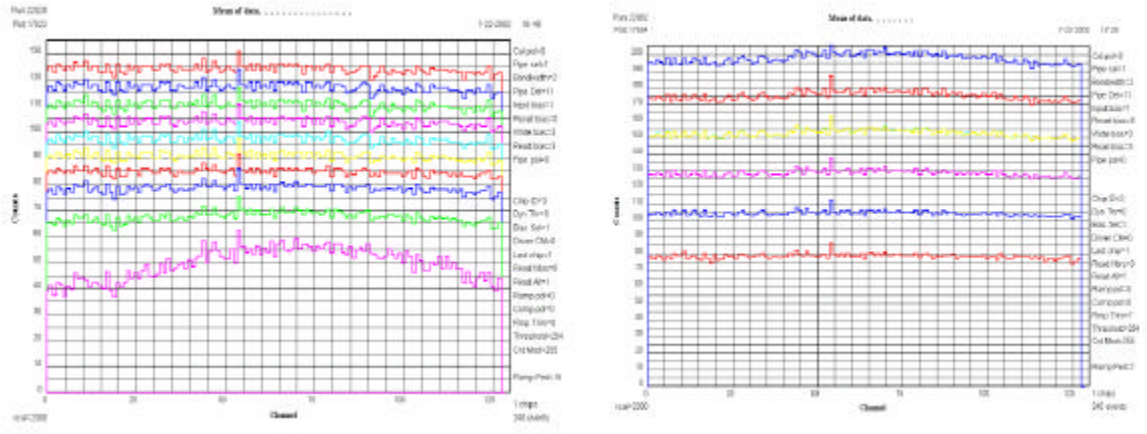

Figure 6 The left histogram shows a pedestal scan from the SVX4. The right histogram shows a ramp range scan. The pedestal scan only includes settings from 0-9. Higher settings give pedestals of zero.

6.1 Pedestals as a Function of Digitization Frequency

The SVX4 chip is extremely fast when digitizing. During digitization, any back-end clock edge seen by the chip will count as an ADC count. The pedestal value is downloaded to the chip, but this value depends on the frequency of the back-end clock. In order to compare with simulation and to see how fast the chip can actually digitize, the digitization clock frequency was increased as high as possible while locking the readout frequency. The results are shown in Figure 7.

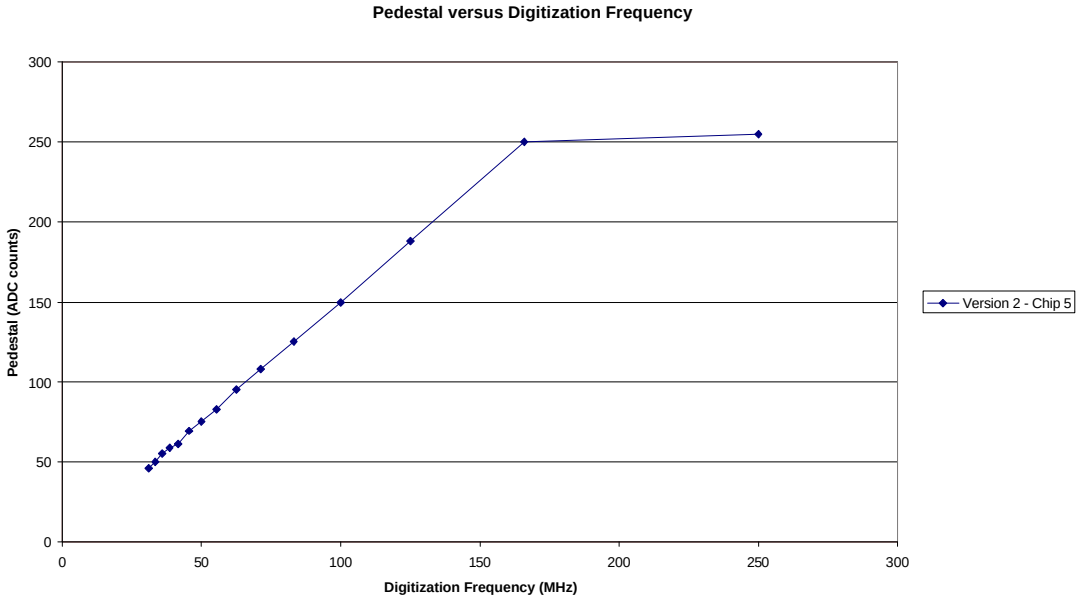

Figure 7 Pedestals as a function of the digitization frequency for a Version 2 chip. The readout frequency of the chip is locked at 25 MHz. The linear increase is expected from theory. As the clock gets faster it hits the counter modulo before the ramp equals the pipeline output. This is evident for frequencies above 150 MHz.

6.2 Variations from chip to chip

We were also interested in studying how much the settings for the pedestal and the ramp varied from chip to chip in order to study process variations. We had five individual SVX4 chips mounted on chip carriers. We performed a ramp range scan and a pedestal scan to measure the variations between chips. In Figure 8, the value of the pedestal is shown for different values of the ramp range that were downloaded. The structure seen is due to the weighting scheme used in adjusting the slope of the ramp. The formula was given in Section 6.

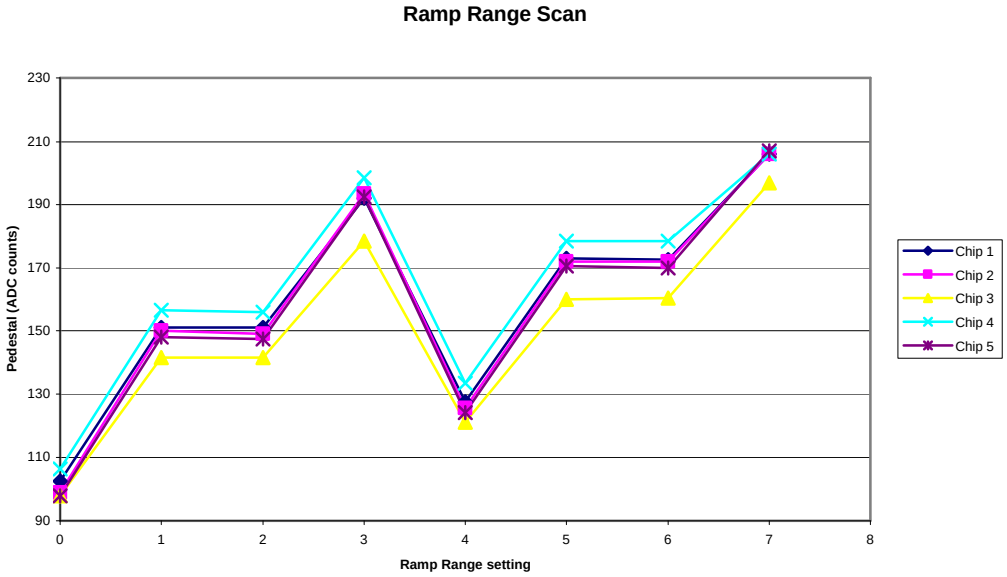

Figure 8 Ramp range scan showing the value of the pedestal as a function of the value downloaded for the ramp. This plot shows that even though ramp range is a three bit allowing 7 values, two values give the same pedestal therefore only giving five different values for the slope for the chosen external resistor on the chip carrier.

In Figure 9, we converted Figure 8 into a plot of the pedestal as a linear function of the ramp slope. This shows the linear nature of the pedestal for different values of the slope. We observed up to a 7.5% difference from chip to chip.
We also measured the variation of the pedestal as a function of the ramp ped setting. This is shown in Figure 10. We can see the linearity up to ramp ped setting of 7. The non-linearity is due to the pedestal bowing[3]. We observed up to an 11% variation for the pedestal settings. Both of these studies were done with a fixed digitization and read out frequency.

[3] We discuss the pedestal bowing in Section 16.1

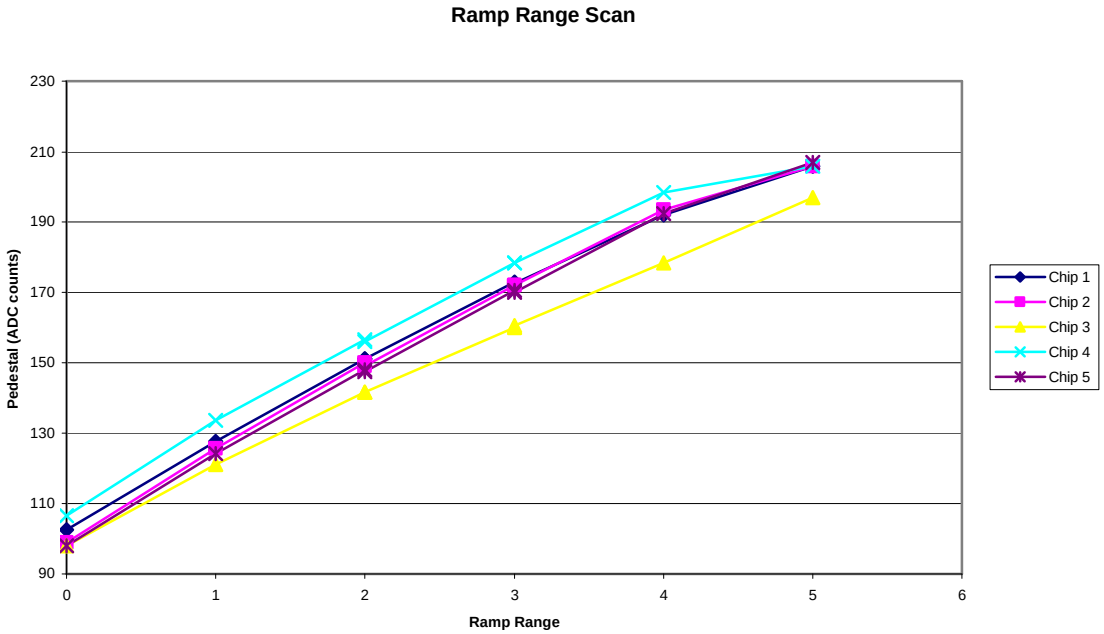

Figure 9 Pedestals as a function of the ramp range slope. The conversion is implicit in the histogram. This shows the linear nature of the slope versus the downloaded value of the ramp range value.

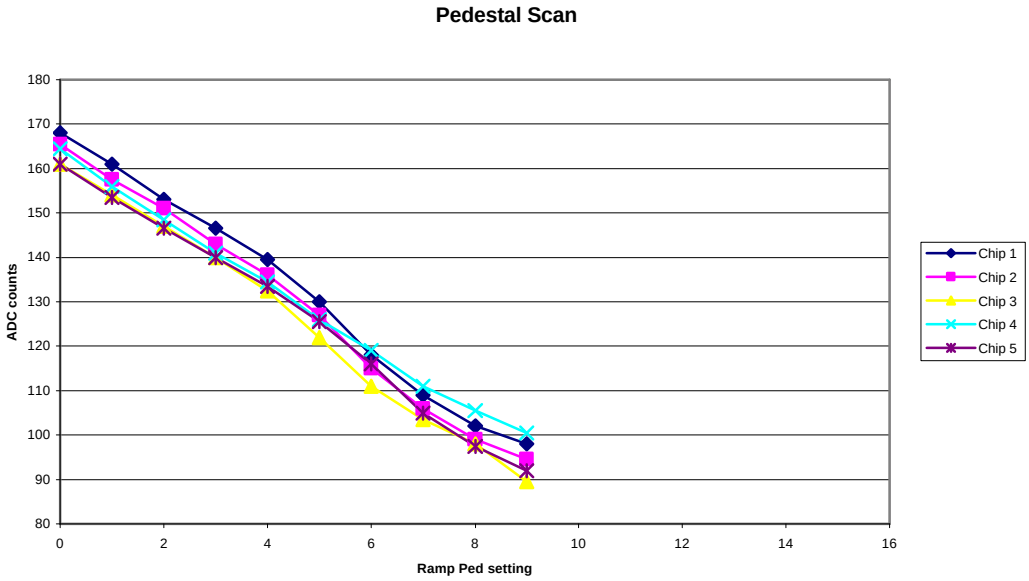

Figure 10 Pedestals as a function of the ramp ped setting. Ramp ped is the variable that downloads the pedestal. The values for the pedestals are slightly non-linear above a setting near 7. This is due to the pedestal bowing discussed in the text. Values higher than 9 give a value of zero for the pedestal.

7 Linearity

We studied the linear response of the channels by injecting a known amount of charge in increments and then monitoring the ADC output. Figure 11 shows the ADC output of a five individual channels of the SVX4 as a function of the calibration injection voltage. In order to measure the non-linearity, we fitted this data to a straight line and then plotted the deviation from the predicted value. We show the results of this process in Figure 12. We see from Figure 12 that the greatest deviation was on the order of 1% and that occurred when the injection calibration voltage was zero.

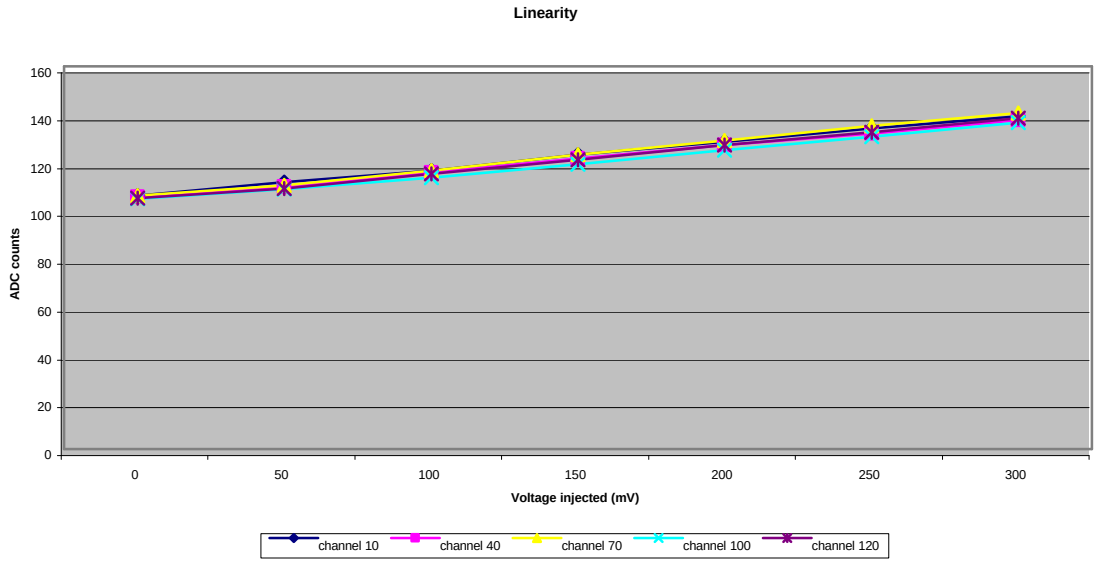

Figure 11 ADC output as a function of the calibration input voltage for five independent channels of the SVX4 chip.

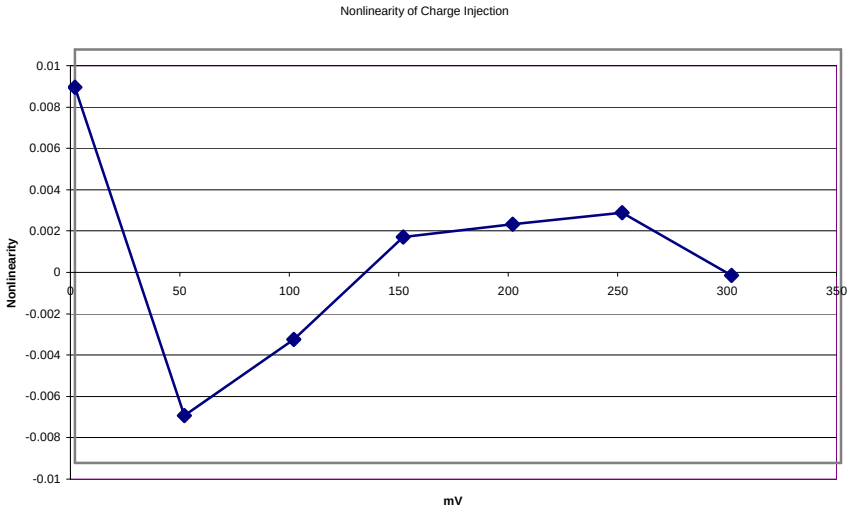

Figure 12 The non-linearity of the ADC output as a function of the calibration injection voltage for a single channel. The largest deviation from the straight line fit occurs when there is no charge injection. It is of the order of 1%.

8 Current Consumption

Because the Stimulus system uses single chips on a chip carrier, we were able to measure the current consumption of the SVX4 using an AM 503 Current Probe Amplifier. The current probe uses the Hall Effect to measure the magnetic field of the cables that carry current. The current probe was calibrated using a 5V supply voltage and a 200 Ω resistor. Expected and measured currents were 25 mA and 24 mA, respectively. To measure to currents separately or summed, a single cable was attached to the power supply with the least ripple, then subsequently teed and connected separately to AVDD and (SVDD+DVDD).

Only one Version 2 chip was measured in CDF mode. We measured the currents for AVDD and DVDD separately and then combined. Also, the SVX4 currents were measured for a variety of chip settings and bias voltages. There was a small dependence on chip output driver current and a dramatic dependence on chip bias. The latter will have a major impact on requirements for low-voltage power supplies.

Data was acquired using single-event read all runs in CDF mode with a calibration pulse injected on every 10^{th} strip. Separately measured currents agreed with their sum to within a few percent. Currents measured in different runs also agreed to within a few percent. Because we used an unsaturated cal inject on every 10^{th} line, SVDD current may not have been maximized during readout.

Figure 13 shows AVDD and (DVDD+SVDD) currents measured separately as well as the sum. Bias at the chip carrier was 2.44V. A green line found in all figures is the trigger signal that was used the trigger the oscilloscope.

In Figure 14, we see the shape of AVDD as a function of time. The jump in the current is when the chip entered the acquire cycle. The shape of AVDD is believed to be from the following. When then chip is not acquiring data and neither the preamp reset or the front-end clock are not held high, the chip enters a saturation state as demonstrated by the long tail in AVDD. It is believed that the current would stay constant if before and after the data acquire cycle, the preamp reset and the front-end clock were held high.

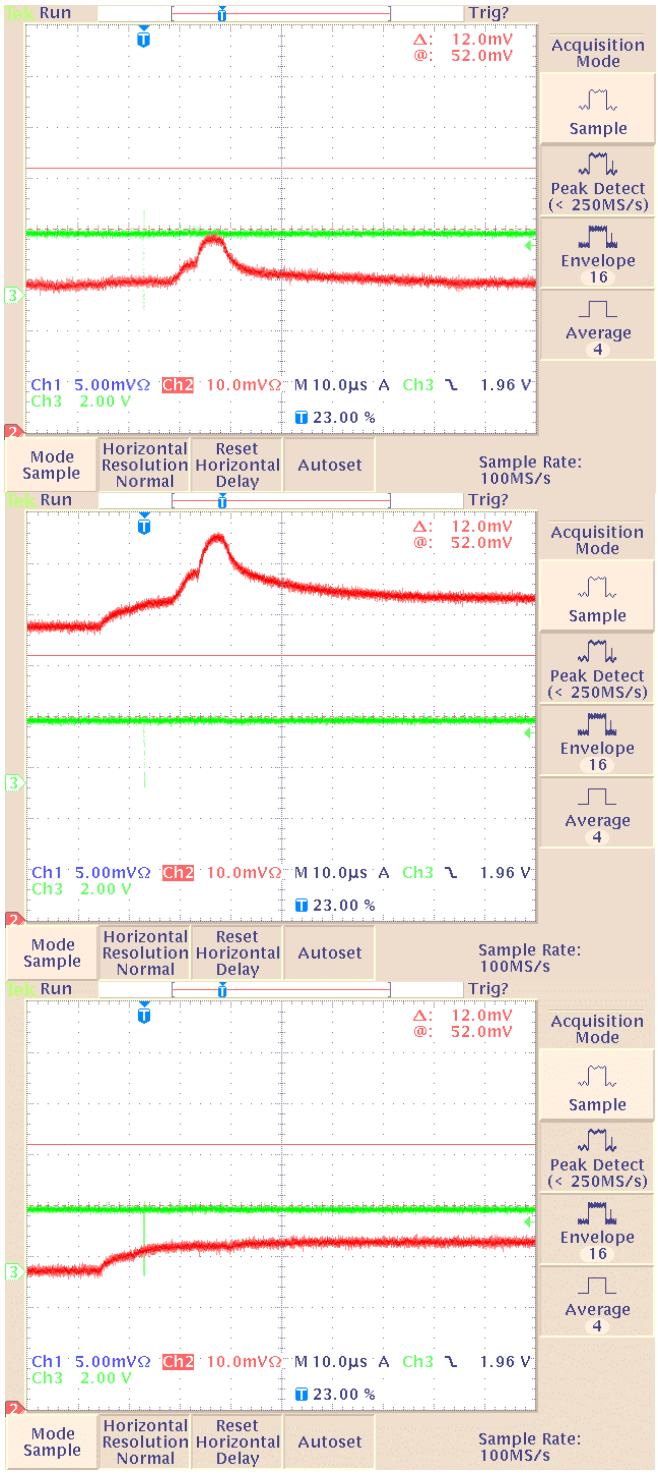

Figure 13 Measured currents from AVDD, (DVDD+SVDD) and the combined voltage supplies are shown on the top, middle, and bottom, respectively. The vertical scale is 20 mA/division and the zero is at the lower left-hand corner. The peak on (DVDD+SVDD) is probably due to readout. Individual currents add to the total current with a few percent.

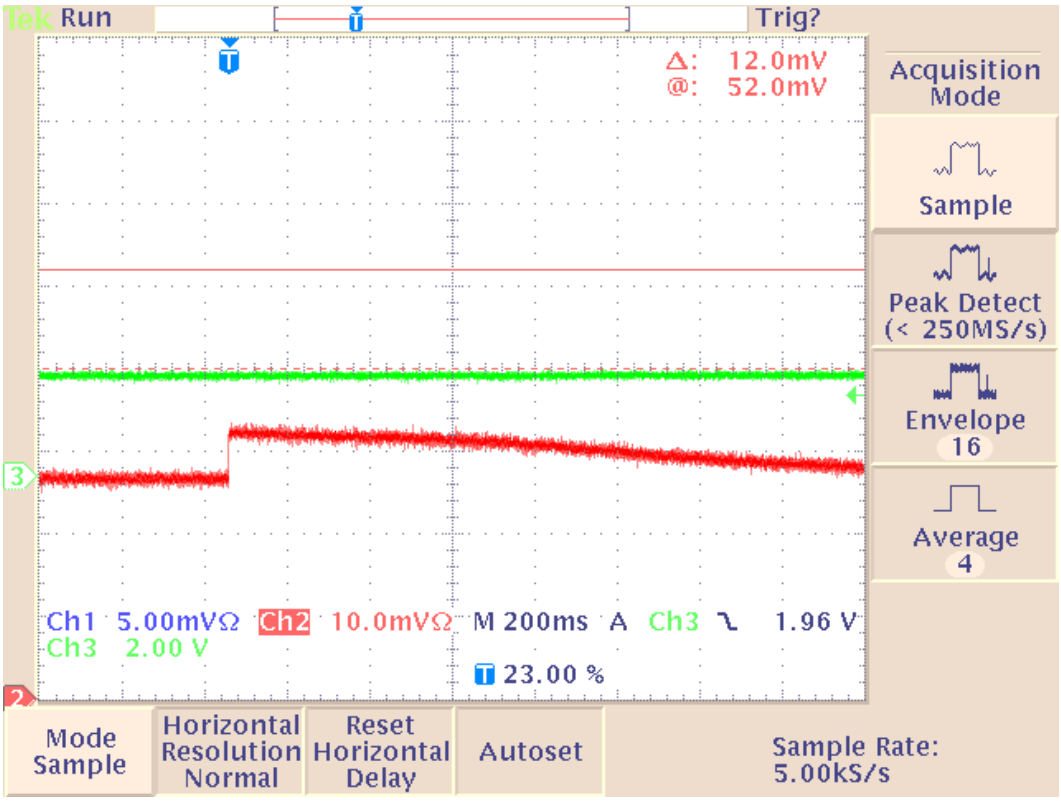

Figure 14 AVDD versus time. The vertical scale is 20 mA/division, with the zero indicated by the red arrow at the lower left-hand corner. The horizontal scale is 200 ms/division. The explanation of the shape is discussed in the text.

The output driver was varied over its maximum range from id=1 (smallest current) to id=7 (largest current). The AVDD current was constant, but the (DVDD+SVDD) current varied by almost 20% as shown in Figure 15.

Finally, total SVX4 current as a function of bias voltage is displayed in Figure 16. The current increases approximately by a factor of two between bias voltages of 2.25V and 2.75V. The bias was measured at the chip carrier; supply voltages were higher by a few hundreds of a Volt.

Though more investigation is required, hybrid currents required probably consist of the peak above wings added to the total current after readout multiplied by the number of chips. For a 10-chip hybrid, that would be (126 mA * 10) + 24 mA = 1.5 A. To test this interpretation, the currents will be measured on a 2-chip hybrid.

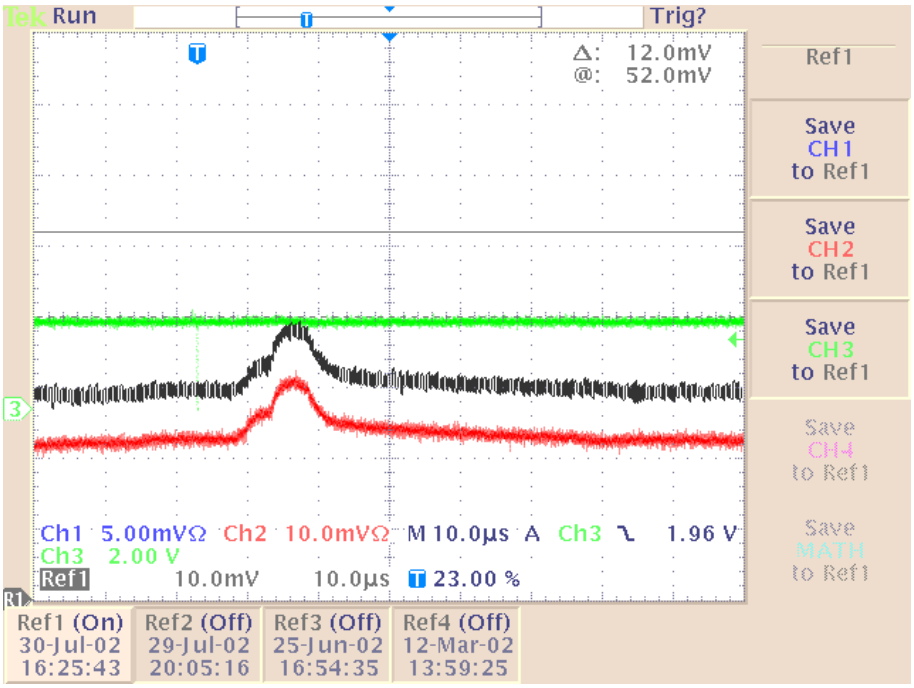

Figure 15 DVDD currents measured as a function of time at 10 ms/division. The higher (black) curve is for id=7 (maximum output driver current), while the lower (red) curve is for id=1 (minimum output driver current). One major vertical division corresponded to 20 mA, and zero was at the lower left-hand corner. Peak currents were 76 and 62 mA, respectively.

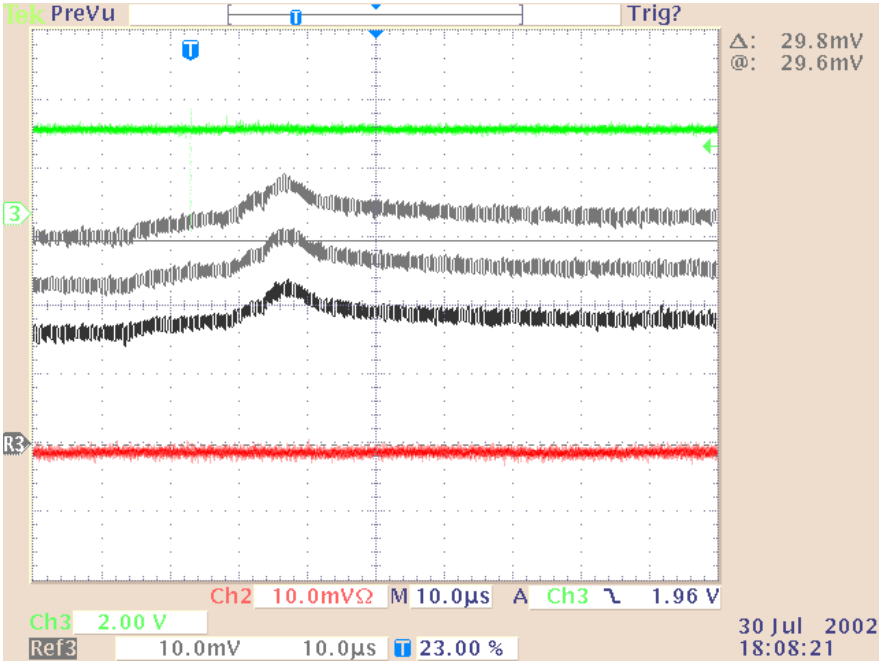

Figure 16 Total SVX4 current as a function of bias voltage with the maximum output driver current (id=7). The highest curve is for 2.75 V on the chip carrier, the middle curve is for 2.50 V and the lowest curve for 2.25 V. The horizontal scale is 10 ms /division, the vertical scale 50 mA/division, and the red line represents zero current.

9 Real Time Pedestal Subtraction

Many effects that contribute to noise (e.g. electromagnetic pick-up) influence all channels on a chip collectively as common mode noise. The SVX4 ADC is capable of performing a common mode (pedestal) subtraction. The subtraction of the pedestal in real time is also known as Dynamical Pedestal Subtraction (DPS). The circuitry for this is shown in Figure 17. For a more detailed study of the DPS circuitry in the SVX3 consult [4].

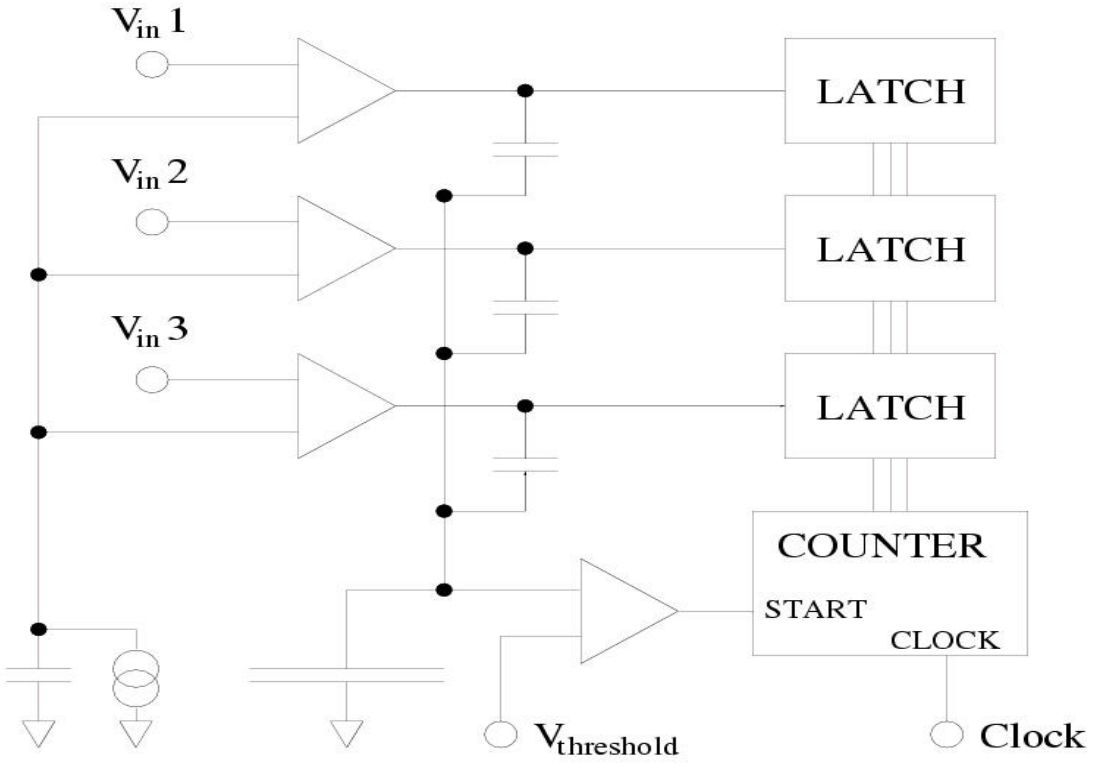

Figure 17 Cartoon schematic of the DPS circuitry inside the SVX4. As the ramp reaches the pedestal value for a particular channel, the comparator fires and increases the voltage on the DPS comparator shown in the middle of the figure. When the DPS comparator fires the clock begins to count, so if many channels are at pedestal the counter begins near the pedestal level.

In order to study the noise contribution in DPS mode, the ADC has an analog delay built in so the baseline of the DPS mode output is not at zero[5]. The threshold for the number of channels needed at pedestal in order for the DPS mode to operate properly is set using an external resistor. The DPS comparator has a 20 kΩ resistor to ground as a default, so the overall or equivalent resistance is calculated by computing the resistance of the default resistance with the external resistor on the hybrid. The equivalent resistance is used to bias a series of current mirrors inside the chip that effectively sets the threshold voltage. In Figure 18, we see the numbers of channels that are needed at pedestal as a function of the equivalent resistance in order for the DPS circuitry to operate properly.

[4] *Statistical Study of SVX3D Chip Dynamic Pedestal Subtraction Threshold Level*, Martin Rehn, Tony Affolder, Maruice Garcia-Sciveres, Igor Volobouev.
[5] Variations of the baseline in DPS mode give a first-order measurement of the noise.

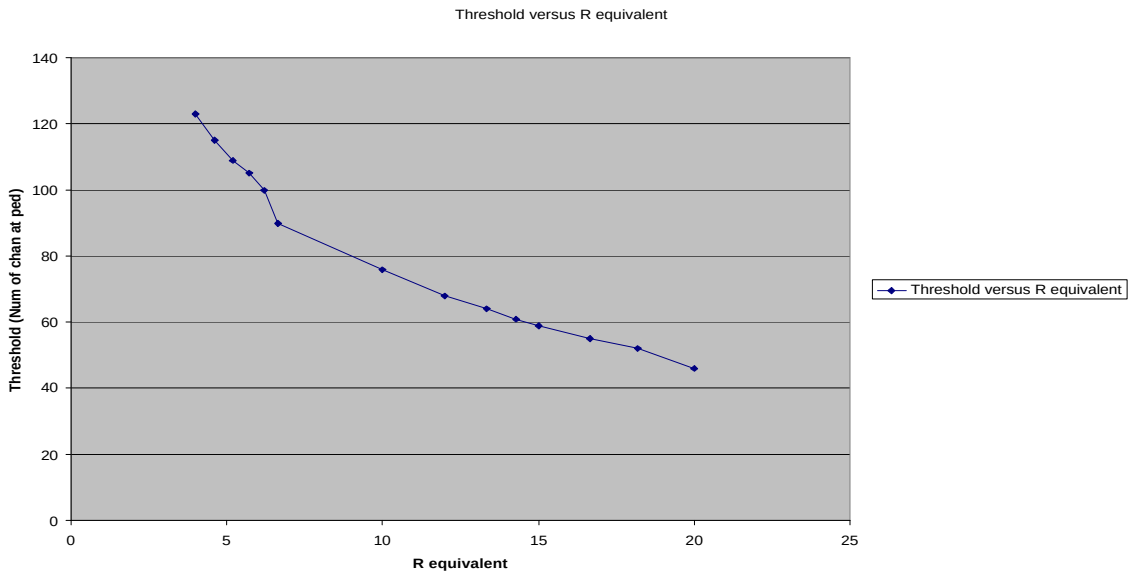

Figure 18 Number of channels needed at pedestal (DPS threshold) as a function of the equivalent resistance of the SVX4 chip in DPS mode. With no external resistor, the default resistance is 20 kW **and the DPS circuitry requires approximately 44 channels at pedestal.**

In order to confirm that the DPS circuitry was working properly, we also measured the relative gain between DPS mode and non-DPS mode for a single channel of the chip as a function of the DPS threshold. In Figure 19, we see the relative gain for a single channel between DPS mode and non-DPS mode of a single channel with a 10 pF load. From the figure, it is clear the gain with DPS mode on and the gain with DPS mode off are similar. In Figure 20, we see the analogous measurement of the relative noise between DPS mode and non-DPS mode. Again, the measurements are taken using the same channel with a 10 pF load. The graph clearly shows the noise between the two modes is similar.

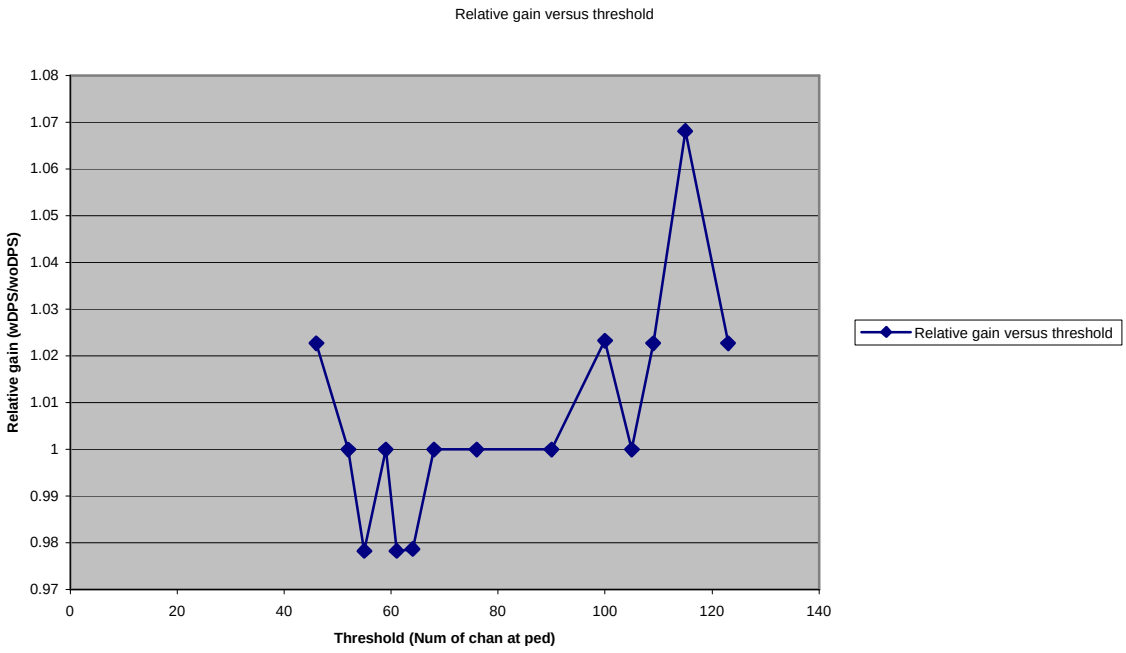

Figure 19 The relative gain of a bonded channel with a 10 pF load in DPS mode to non-DPS mode as a function of the DPS threshold. It clearly shows that the gain is similiar between the two modes.

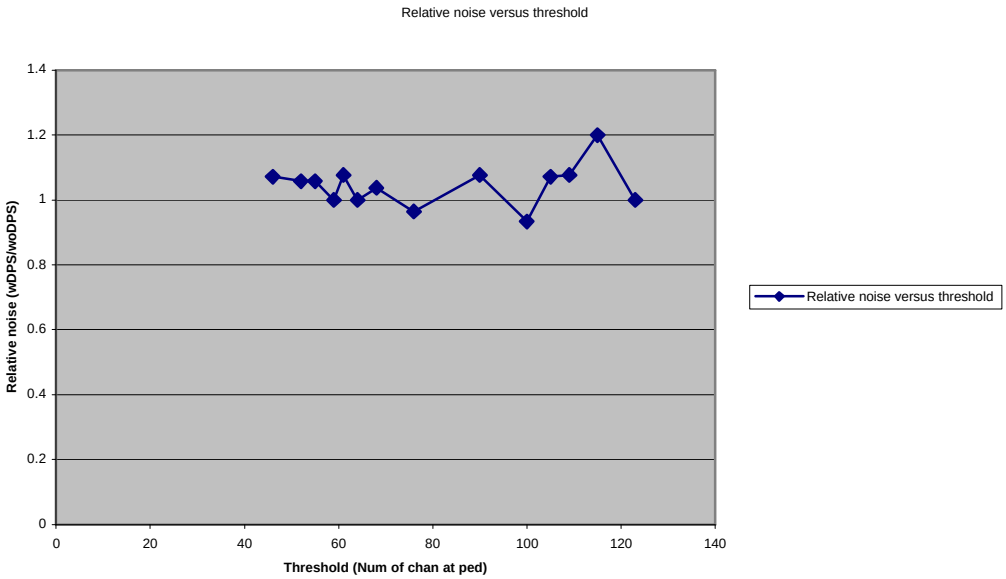

Figure 20 The relative noise of a bonded channel with a 10 pF load in DPS mode to non-DPS mode as a function of the DPS threshold. It shows that the DPS mode does not introduce more noise into the system.

10 Readout

Unlike the precursors to the SVX4, it does not use a collapsible output FIFO, but instead utilizes a token system. To confirm the functionality and to compare against simulation, we made the following measurements.

10.1 Output Driver and Current

We measured the output driver current of the data lines by placing scope probes on both sides of the differential bus line terminated by a 100 Ω resistor. By measuring the voltage on each side of the bus, the current can then be calculated. Table 7 shows the measured values as a function of the control bits downloaded to the chip.

Bit	Measured	Design
001	5.6 mA	5.6 mA
010	9.4 mA	9.2 mA
100	13.2 mA	13.4 mA

Table 7 Output currents for the data lines from the SVX4 chip as a function of the downloaded bits. The measured values agree well with the design values.

10.2 Double Readout

It is known from simulation that the readout of the SVX4 chip in sparse mode will produce a double readout of a single channel in certain conditions because of the token system. We are able to simulate the conditions for double readout using the following procedure. We set the readout clock frequency and then injected on the highest channel number. If we had a double readout, then we would inject on a lower channel number until the double readout disappeared. This way we could calculate the propagation time of the token (we know the clock frequency) and then predict where the double readout would occur for different readout clock frequencies. This worked very well. In Figure 21, we show the threshold for double readout as a function of the readout frequency. A clear linear relationship is observed. As the readout frequency gets higher, the threshold for double readout is lowered.

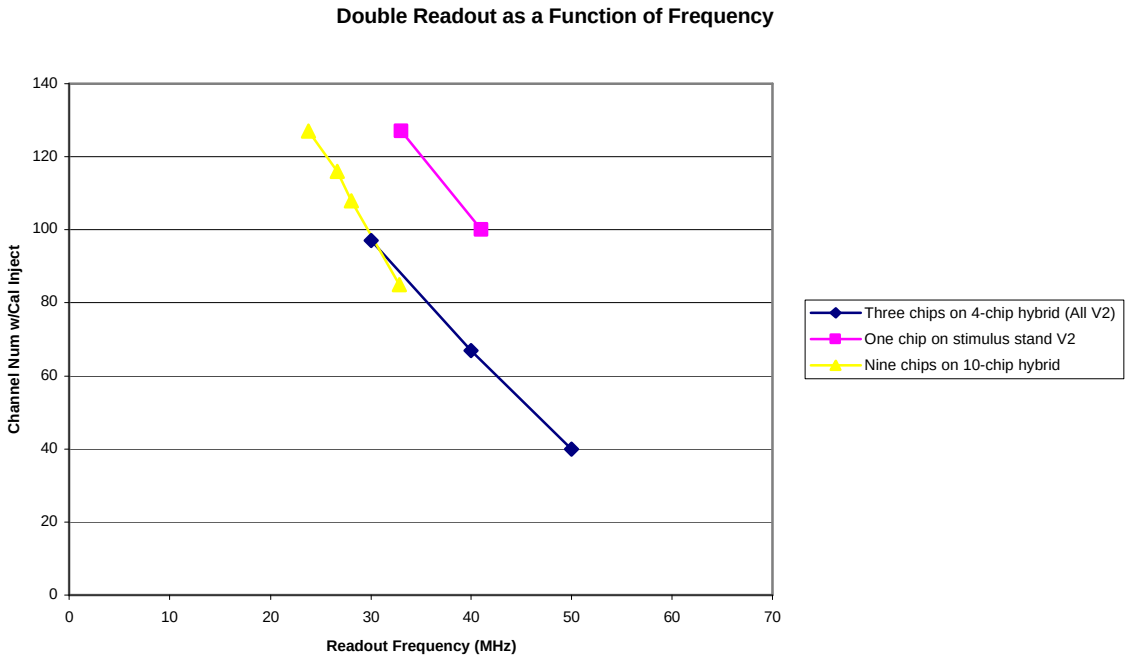

Figure 21 Threshold for double readout as a function of the readout frequency clock. For example, for three chips on the four chip hybrid, setting the readout clock frequency to 40 MHz will give a double readout if the charge is injected on channel 62.

It was found that no double readout occurred for the first chip on the four or ten chip hybrid during this procedure. Through further investigation, a double readout could be produced if the time between the signal Priority In going low and the first back end clock pulse was less than one-half the readout clock frequency. This is shown in Figure 22. If the time between Priority In and the first back end clock was less than 19 ns, a double readout of the first chip occurred.

We also conducted a study to observe how the threshold for double readout depended on the DVDD voltage supply. From simulation we know that the speed of the token is increased as the DVDD voltage supply is increased. In Figure 23, we show how the threshold for double readout increases linearly with increasing DVDD voltage. We used four different chips on individual chip carriers for this study. It is important to note that at normal frequencies (operating frequencies in the experiment) we could not produce a double readout for single chips. Therefore in this study we increased the frequency until a double readout occurred for the individual chips. It is currently not understood why single chips only produce a double readout with higher than normal frequencies. It is believed to be due to the way that the chip is currently biased on the chip carrier.

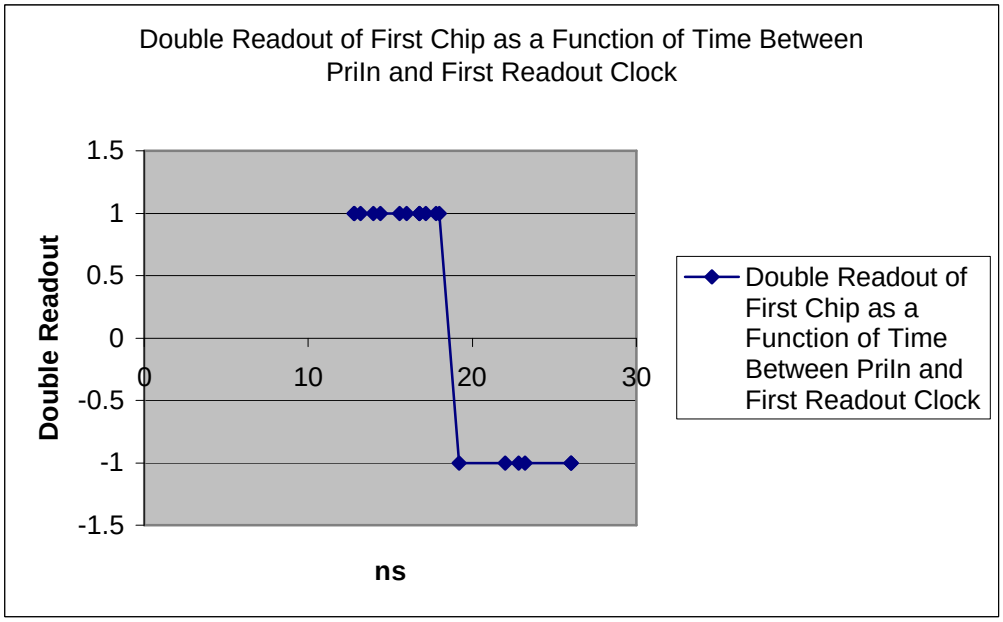

Figure 22 Double readout of the first chip of the ten chip hybrid as a function of the time between Priority In going low and the first readout clock edge going hi. The value 1 is used in a Boolean fashion. This means a double readout occurs when the value is 1 and the double readout does not occur when the value is –1. The transition occurs at approximately at 19 ns which is one-half of a clock cycle at 26.5 MHz.

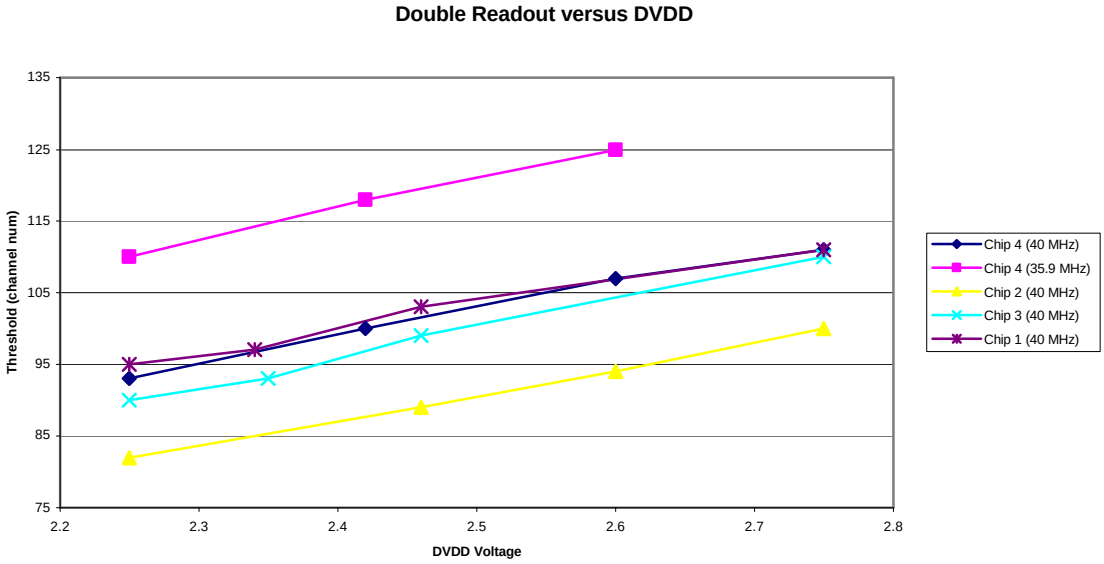

Figure 23 Threshold of double readout as a function of DVDD. It is clear from the figure that as DVDD is increased, the threshold for double readout increased because the speed of the token is increased. These results were obtained using four different single chip carriers at higher that normal back end clock frequencies.

We also studied the double readout using different configurations of the charge injection mask. In Table 8, we show only the data output from an SVX4 chip when the double readout conditions are satisfied. In the far left column of the table, only one channel has charge injected into it. It should be noted that the sparsification feature of the chip is turned on and the threshold for readout is placed high enough to ensure only those channels with charge injection will be readout. The second column shows that if multiple channels are injected on, then only the first channel of the pair gives the double readout. The third column confirms that if the charge injection occurs much lower than the threshold that no double readout occurs. Other columns are obvious, but the last column shows the Read Channel 63 feature of the chip prevents double readout from occurring.

71	7d	00	00	00	00	00
71	7d		7f	7e	01	01
	7e		7f	7e	7e	3f
				7f	7e	7e
					7f	7f

Table 8 Output from the SVX4 showing the double readout for various configurations. The final column shows that turning on the feature to Read Channel 63 all the time prevents the double readout from occurring.

10.3 Additional Readout Features

The SVX4 has additional readout features whose functionality was confirmed by all test stands. For example, the Read All feature was confirmed for non-DPS mode and DPS mode. The Read Neighbor bit was confirmed to work as well. Read Channel 1, Read Channel 63, and Read Channel 128 were all confirmed to be functional. These features were added for different reasons. Counter Modulo and the Digital Threshold were also confirmed to be functional.

10.4 Frequency Margin

In experimental conditions, it is impossible to transmit a perfect clock to the components of the system due to transmission through long cables. So, the chip must be able to function or be able to readout if the back end clock does not have the optimum frequency (f = 25.6 MHz) or 50% duty cycle. We tested this ability by varying the frequency and duty cycle of the back end clock during digitization and readout.
Worst case conditions require that the chip operate within a 20% frequency margin and a 40/60 -60/40 duty cycle. We have confirmed that the chip works well within the required frequency margin and duty cycle for both digitization and readout. At each frequency and duty cycle, a pedestal run was taken along with a calibration injection on every tenth channel to confirm proper operation of the chip. Also, we found that we could increase the speed of readout up to 50 MHz and the chip would function.[6]

[6] The Stimulus System DAQ is limited to 50 MHz because the FIFOs on the Adaptor Board are limited by that clocking speed or ability to clock in data.

11 Black Holes (Pinholes)

In the silicon detector, the AC coupling capacitor that connects the biased silicon strips to the input of the preamplifier of the chip can be damaged or destroyed. This causes a large amount of current to flow into the preamplifier completely saturating the input of the channel. Because of such a large current flow and the input of the amplifier being forced to ground, it becomes possible to forward bias the diode structure between the input channel and substrate of the chip. Therefore, the input channel with the shorted capacitor shares charge with the neighboring channels. In Figure 24, we see a cartoon schematic showing a shorted coupling capacitor and charge sharing through the substrate.

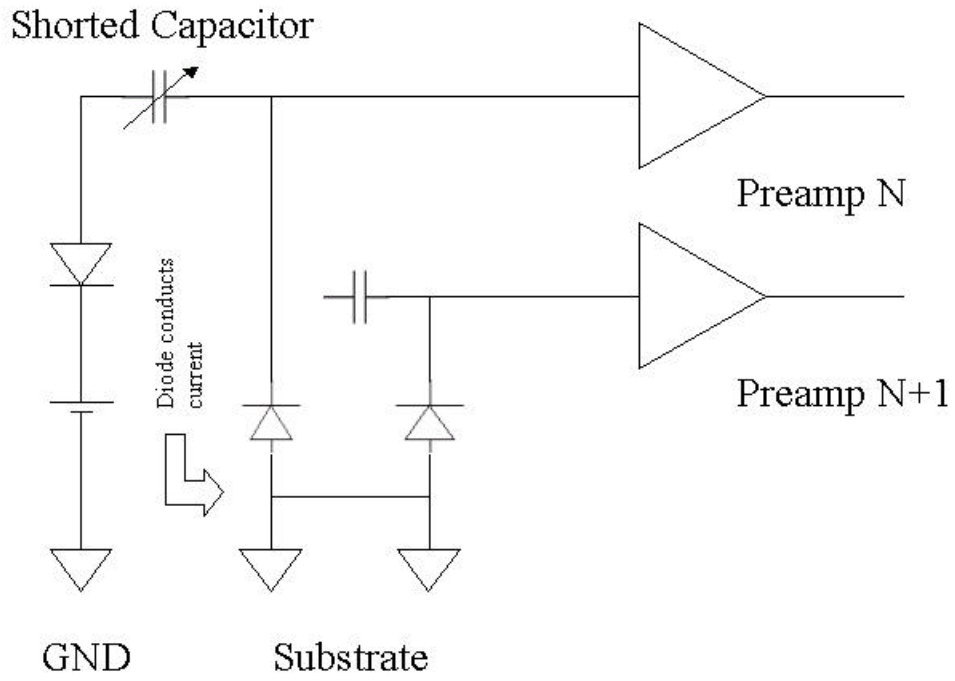

Figure 24 Cartoon schematic showing the charge sharing between neighboring channels when a pinhole exists on the silicon sensor.

When such a situation occurs, the pedestal saturates and the noise drops to zero with the neighboring channels having higher noise. There is no way to disconnect the channel with the shorted capacitor once the detector is installed. Therefore, the input from the channel with the pinhole and the neighbor channels must be rejected and therefore a loss of fiducial coverage and noise is introduced into the system. If one looks at the noise as a function of channel number, one sees the noisy channels around the pinhole and a zero value for the channel with the pinhole. This looks very similar to a space-time diagram of a black hole and therefore the name was adopted.

The SVX4 has a new feature which allows the channel with the black hole to be disconnected so that neighboring channels are unaffected by the large draw of current. In the downloaded bit stream, a bit controls whether the channel mask is used as a charge injection mask or a mask against black holes. Essentially, when the mask is used to disconnect pinholes, it is forcing the preamp reset for the channel to be active during data taking and it is able to source the current from the damaged capacitor.

To simulate a black hole and confirm the new feature of the chip works properly, we did a study using a current source connected directly to the input of the channel (with no coupling capacitor). We increased the current in a linear fashion. We sourced up to 3.5 mA and saw no charge sharing at all between neighboring channels without the black hole mask active. We have also confirmed that the black hole fix does work by activating the mask while injecting a square pulse into the channel and seeing no response. We investigated any possible unknown systematic effects by sinking up to 3.5 mA. In conclusion, it appears that the channels are robust enough not to be affected by a pinhole.

12 Voltage Supply Issues

We investigated the ability of the SVX4 chip to survive over voltaging without incurring damage. For this study, we combined the two power supplies to the chip (AVDD and DVDD) into one power supply to simplify the procedure.

The test consisted of setting the power supply voltage to a set value and then taking a data cycle with the test stand[7]. If the power supply voltage was outside the specifications (2.25-2.75V for both AVDD+DVDD), then after that data run the voltage supply was brought back to 2.75V and a new data cycle was taken and compared to previous data runs. This was done to confirm the proper operation of the chip and to confirm no damage to the SVX4 had occurred.

The point of this test is to explore the voltage ranges that the SVX4 can handle without being permanently damaged. The results were as follows: After 6.8V, we vaporized a wire bond connected to DVDD because the current draw was greater than 1A. From this we conclude that it is more likely the experiment will vaporize a wire bond instead of damaging or destroying the SVX4 chip from an over voltage on the chip.

Another feature of the SVX4 that was investigated was the possibility of latch up when turning on the power supplies to the chip. The SVX2 has a controlled sequence of how the power is applied to the chip to prevent such a situation from occurring. Since the Stimulus test stand has manually controlled power supplies, we randomly turned on the supplies for AVDD and DVDD. We also randomly lowered and raised the power supplies and could not produce any latch up behavior as observed with the SVX2. It is important to note that if the voltages were taken outside the specifications of the chip, it was sometimes necessary to reinitialize the chip to get proper behavior, but no latch behavior was observed.

13 Differential Non-linearity Studies

These studies were done with the Dzero teststand and consists of connecting a sinewave voltage source with a very low frequency. You are trying to measure the difference between the even and odd bits of the ADC. One bit will not give that same amount of charge as the other.

By using a slow sine wave, one can sweep across the entire quantization size of the ADC. The quantization size is $Q=FSR/2^n$. The speed of digitization is much faster than the sine wave so there are no timing issues.

[7] A data cycle consisted of 240 single events.

If the quantization for even and odd values are different, then the following graph would have an even odd effect for the ADC output. To get a measure of the non-linearity, one sums up all the even ADC contributions and then sums up all the odd ADC contributions (for a certain region of the graph) and then uses

(even-odd)/(even+odd)

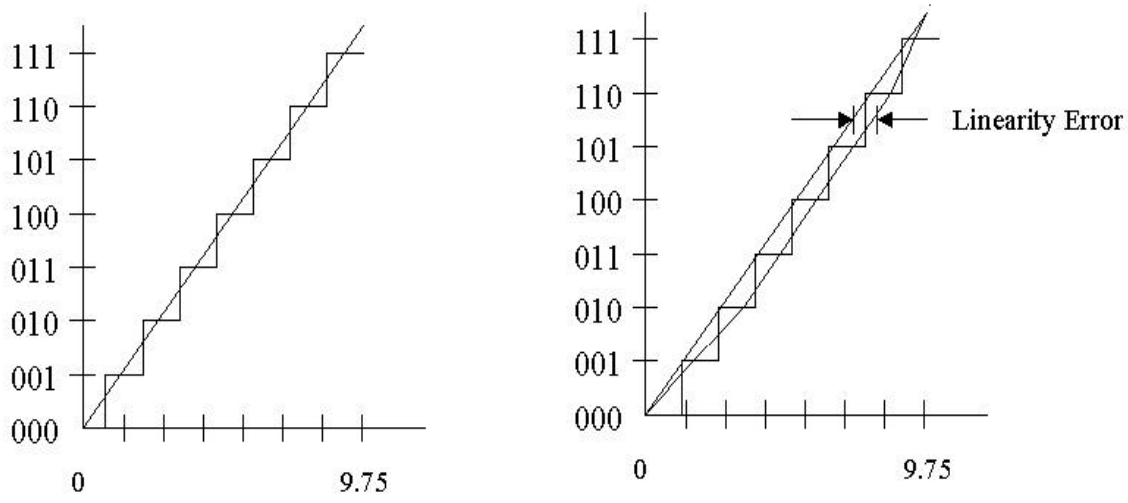

Figure 25 Cartoon representation of the transfer function of a three bit ADC. The left picture shows a perfect conversion, whereas the right hand picture shows an ADC with a slight non-linearity.

We are trying to see if the odd values have a larger integration time, so therefore they will require more charge in order to increment the counter.

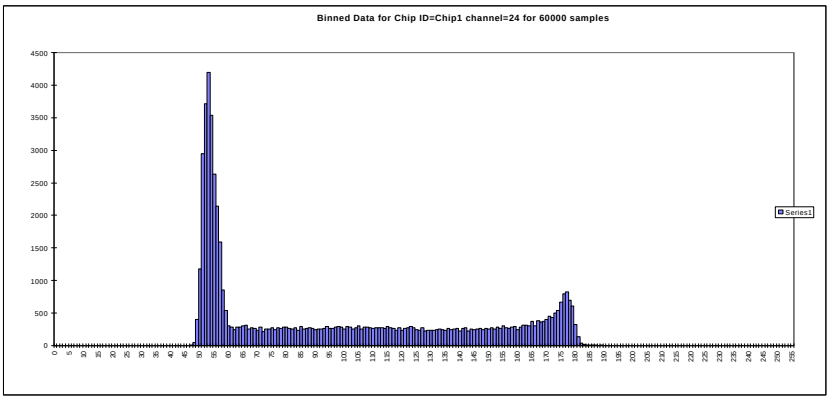

Figure 26 Differential nonlinearity for a ten chip hybrid. The offset is because of the analog delay in the comparator. The plot is not symmetrical due to the voltage divider effect on the calibration input.

14 Data Persistence

We also investigated the data persistence using the D0 hybrids.
The object of this measurement is to see if the resistance of the ends of the capacitors of the pipeline has a high enough resistance to the substrate to prevent substantial leakage of charge of the pipeline capacitor and therefore introducing an error into the measurement of the charge injected into the channel.
The resistance of the MOSFETs that are used as switches on the pipeline have large resistances to start with. By measuring the time of the persistence of the data, one can estimate the resistance of the capacitors ends to the substrate. The value of the capacitance of the pipeline cells is C=10 pF. We measure the time of decay to be 1 ms before the baseline of the pedestal started to shift.

15 Electrostatic Discharge Testing (ESD)

01/26/2001

ESD Test Setup Schematic

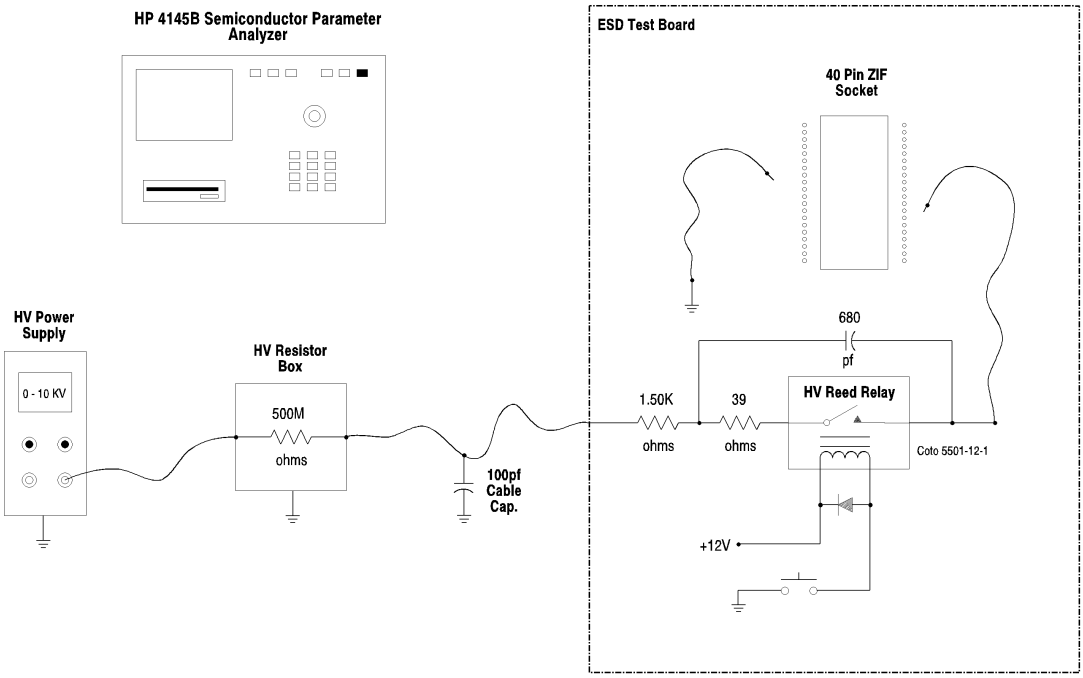

Figure 27 ESD test setup showing the use of the human body model.

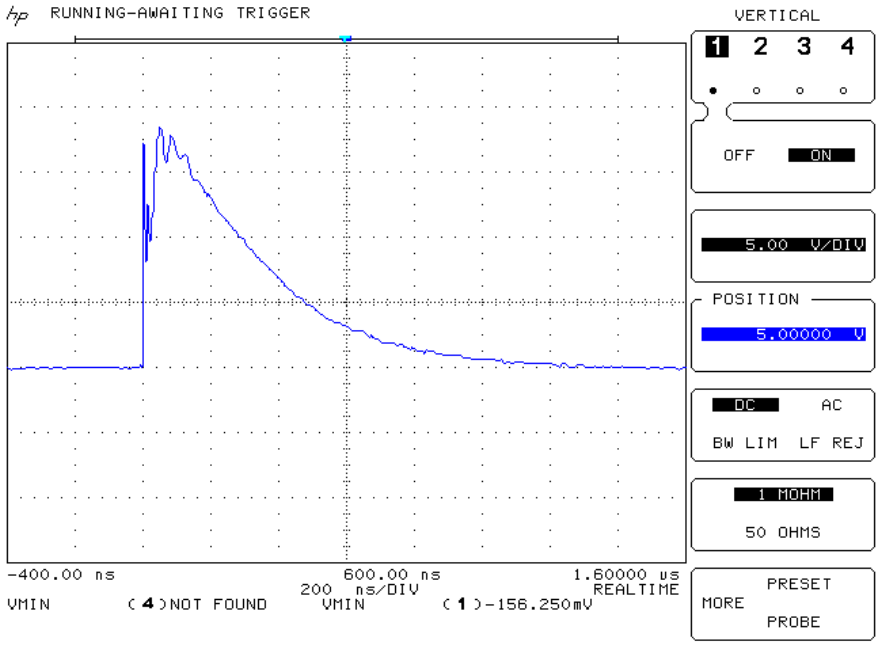

Figure 28 Output waveform from the ESD setup.

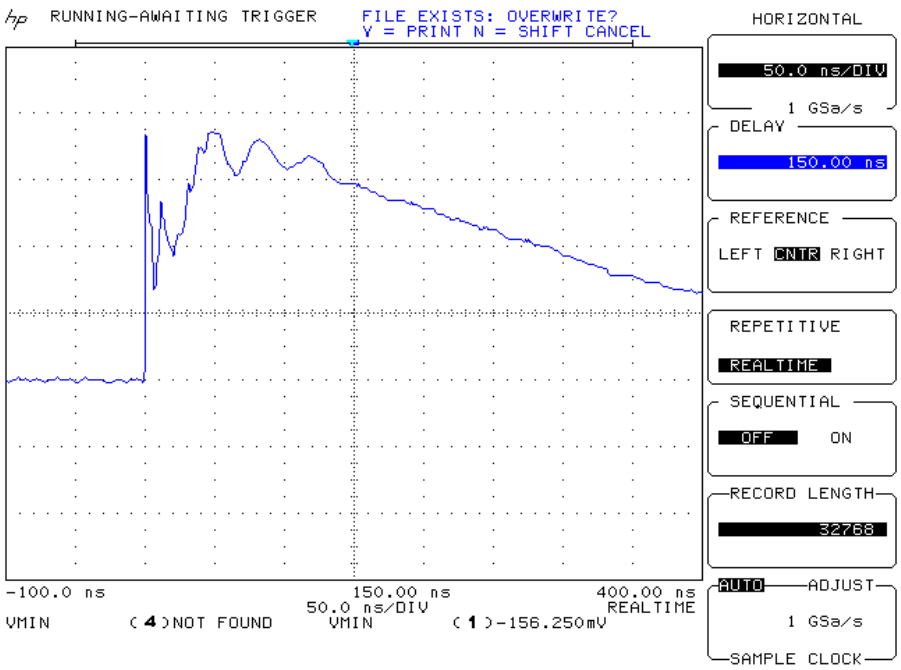

Figure 29 A magnified view of the output waveform from the ESD test setup.

16 Performance Issues

The current SVX4 design currently has three problematic issues: pedestal bowing, channel to channel variations, and pipeline slope. In the following sections we will describe in more detail the studies that have been done.

16.1 Pedestal Bowing

The pedestal bowing is conjectured to be coming from the slew time in the comparator delay. By increasing the I quiescent current, we are able to reduce

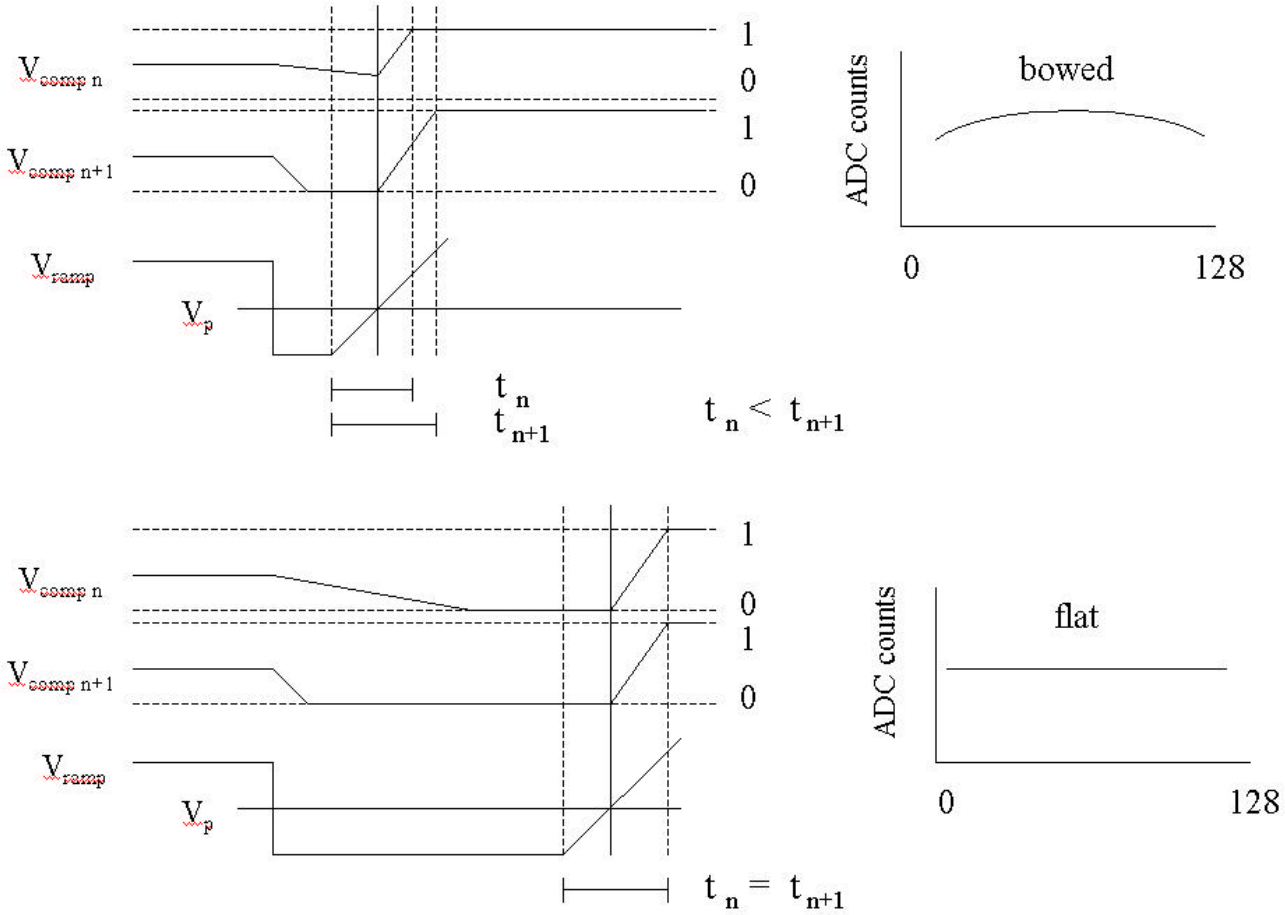

Figure 30 Cartoon showing how the bowing can occur in the SVX4 and how the channel to channel variations happen due to the slew time or speed of the comparators.

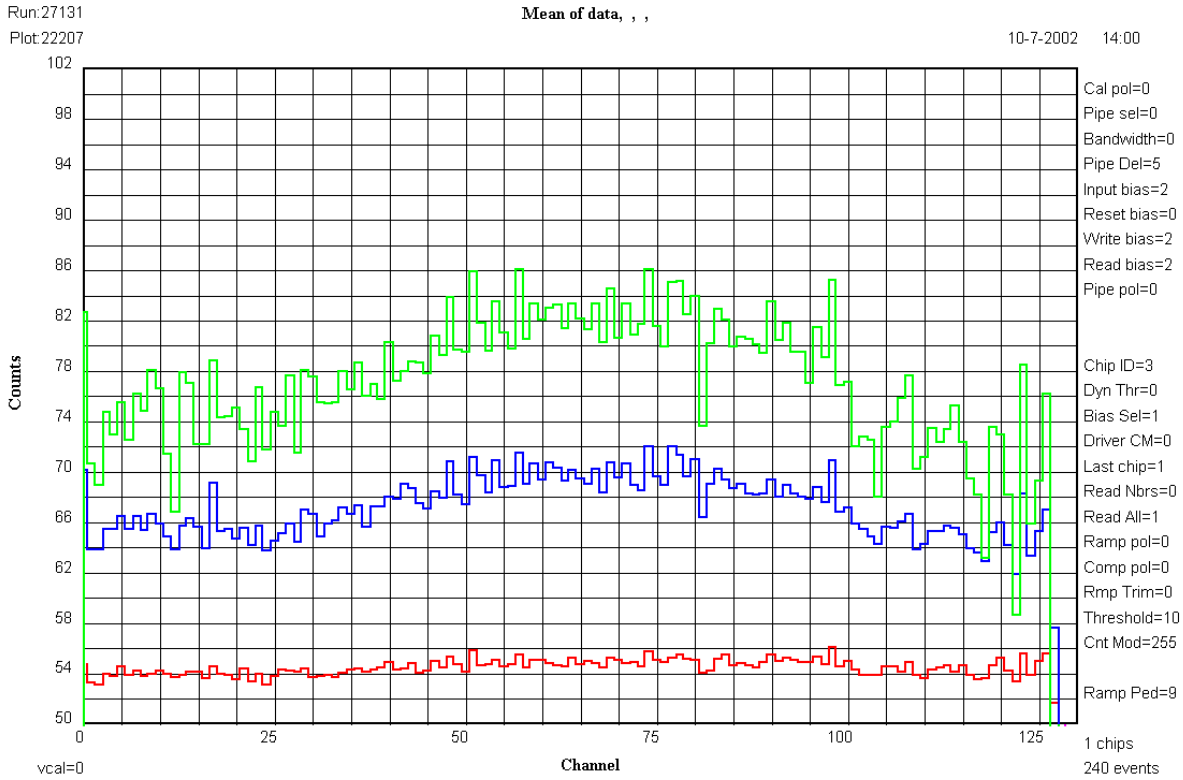

Figure 31 The bowing as a function of the I quiescent current of the chip. As the I quiescent current is increase the bowing decreases. This is because the larger the I quiescent the faster the comparators become. The green line corresponds to V=510 mV, the blue line corresponds to V=530 mV, and the red line corresponds to V=610 mV.

We also investigate whether or not there was a voltage drop along the power bus that supplies voltage to the comparators. We did this by producing a voltage drop along the DVDD bus. We wirebonded one side of DVDD and then connected the other side through a resistor to ground. This produce a noticeable voltage drop along the bus, but we saw no effect in the bow.

16.2 Channel to Channel Variations

We have found a similar effect on the channel to channel variations by increasing I quiescent. The channel to channel variation are still there but the spread is surmised to be a function of the comparator delay.

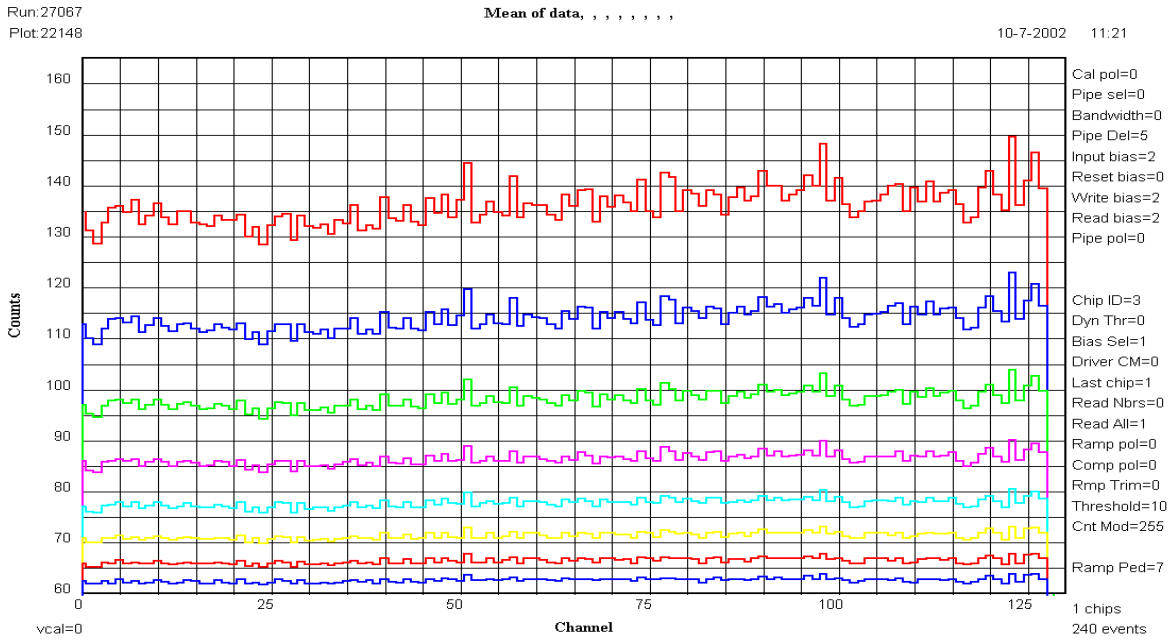

Figure 32 Channel to channel variations as a function of the I quiescent current to the chip. This transparency only contains ramp up values.
Red V=490 mV ped = 135 I slope= 117 uA
Blue V=510 I slope= 153
Green V=530 I slope= 190
Maroon V=550 I slope= 235
Light Blue V=570 I slope= 294
Yellow V=590 I slope= 337
Red V=610 ped =65 I slope= 398
Blue V=630 ped =62 I slope= 461 uA

We also checked whether or not the channel to channel variations are coming from possible charge injection when the reset switch is released. We did this by changing the gain of the chip, or the ramp, and then measuring the channel to channel variations. We saw no difference in the magnitude of the variations as a function of gain.

16.3 Pipeline Slope

The is a systematic increase in the average pedestal as one looks at higher pipeline cell numbers. The leading explanation is that a considerable resistance has developed along a power trace connecting all the preamplifiers in the pipeline. It is believe that this is coming from a resistance that has developed along the power bus the slightly changes the voltage from each pipeline cell.

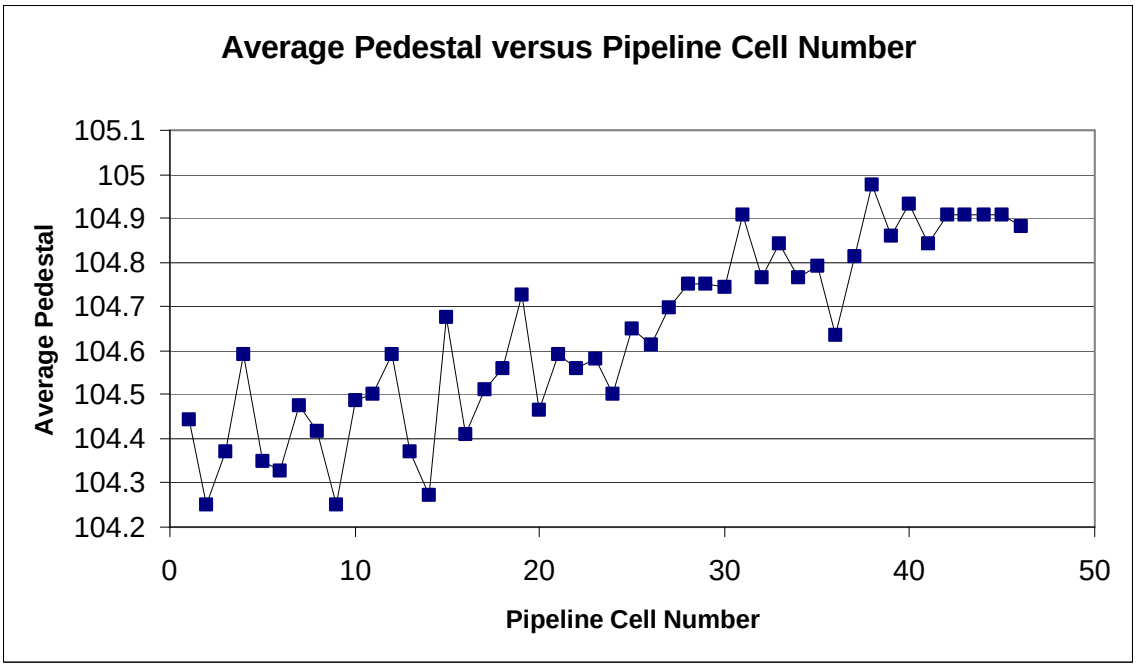

Figure 33 Average pedestal as a function of the pipeline cell number for an individual channel. The channel number was number 63.

In order to confirm its dependence we fit a linear and a quadratic to the slope.

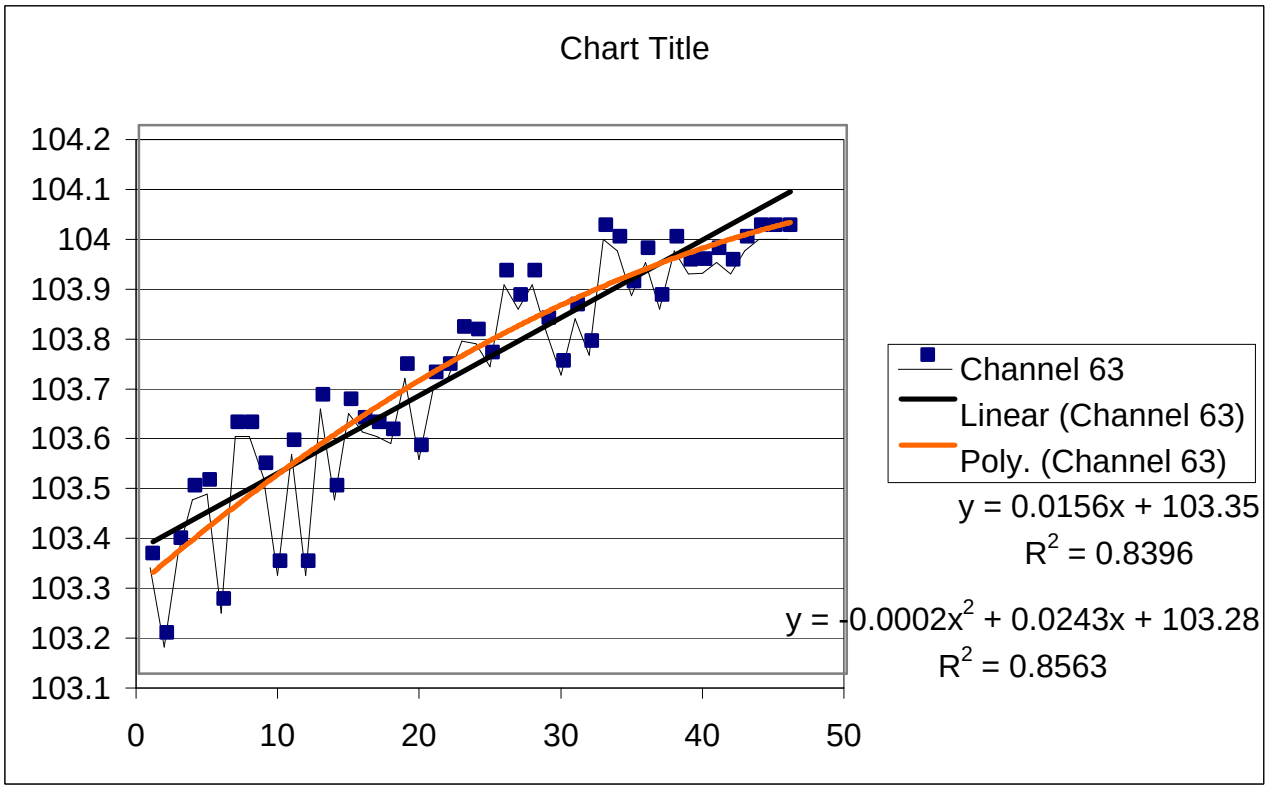

We also carried out a exhausted study of the slope of the pipeline as a function of channel to investigate if the bow was related to the geometry of the chip.

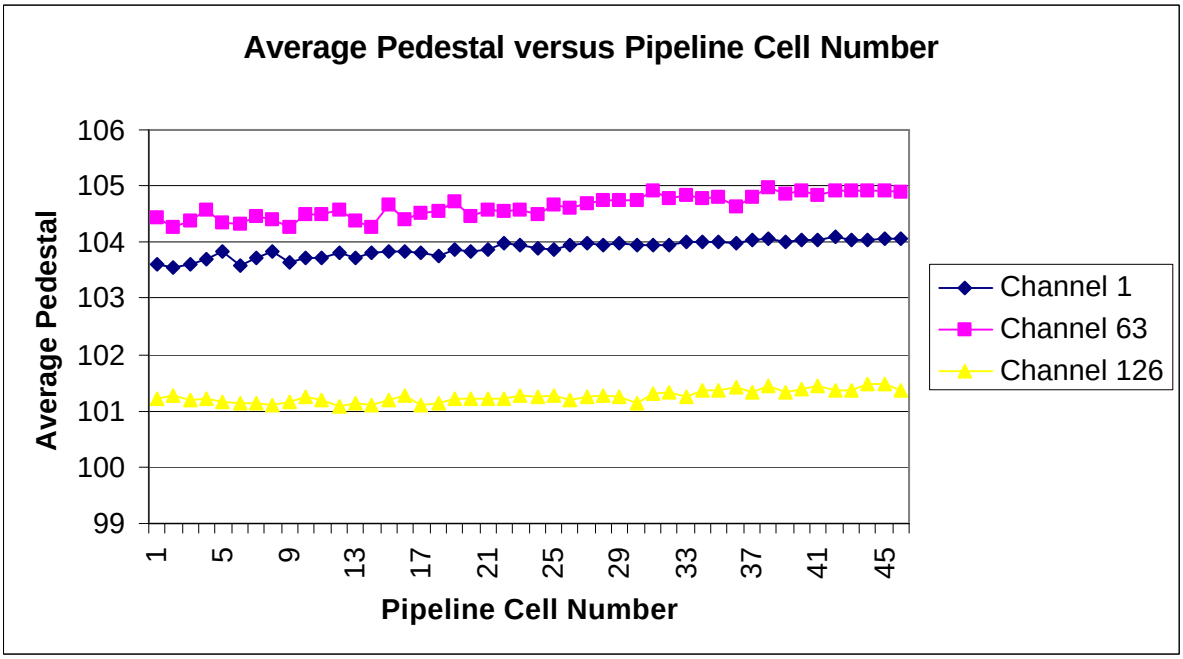

Figure 34Comparison of pipeline slopes for different channel numbers. The difference in pedestal is due to the bowing effect discuss above.

In order to confirm that the slope in the pipeline was indeed coming from something other than digital coupling, we turned on the preamp reset throughout the acquire cycle and took pedestals.

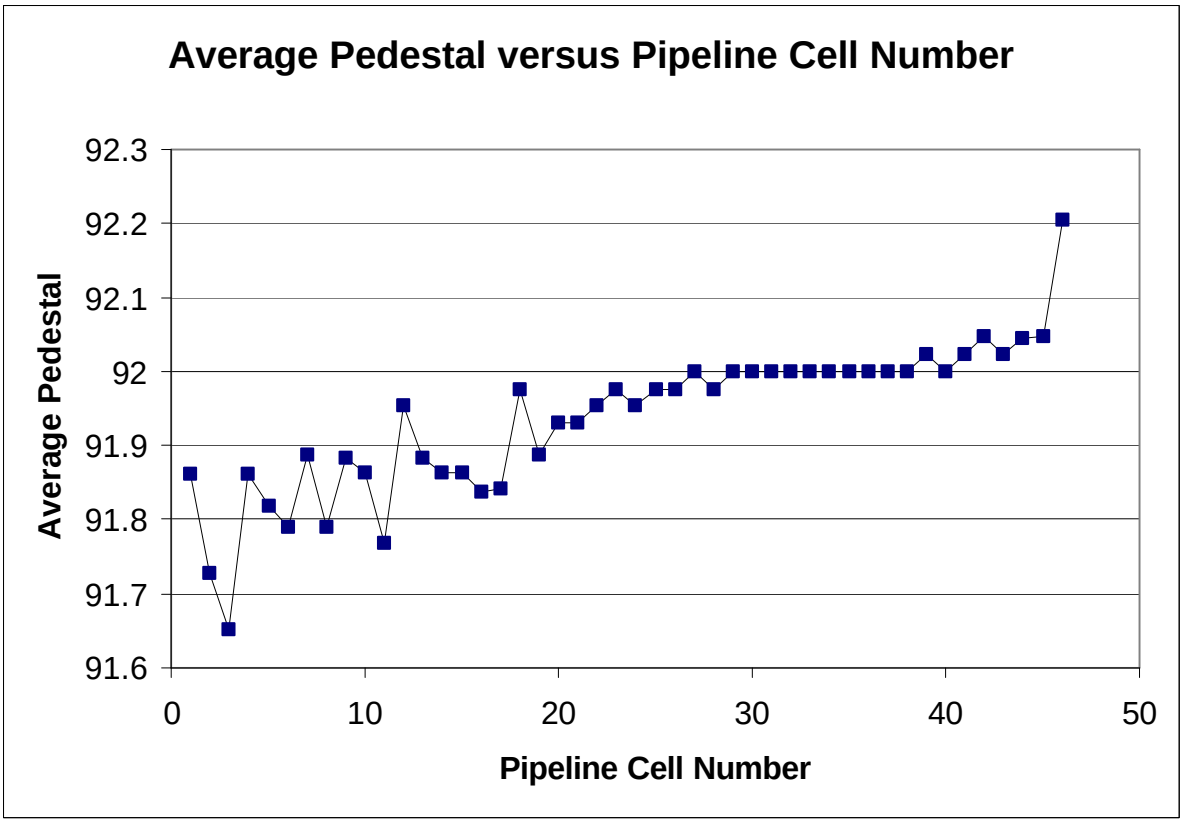

Figure 35 Average pedestal versus pipeline cell number. The data was taken with the preamp reset held high through out the acquire data cycle. It is from channel 63.

We also carried out exhaustive power studies.

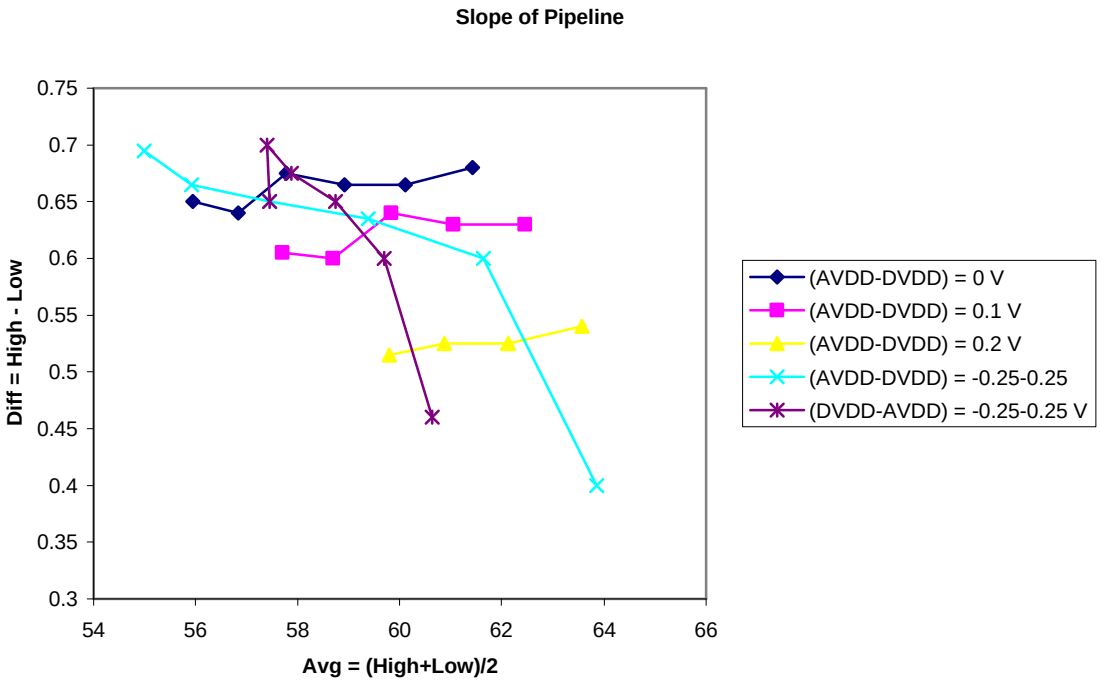

Figure 36 Slope of the pedestal as a function of the voltage supplies.

16.4 Receiver Failure

It was observed that too much I quiescent current or too little and the chip does not work. The reason was found to be a faulty layout.

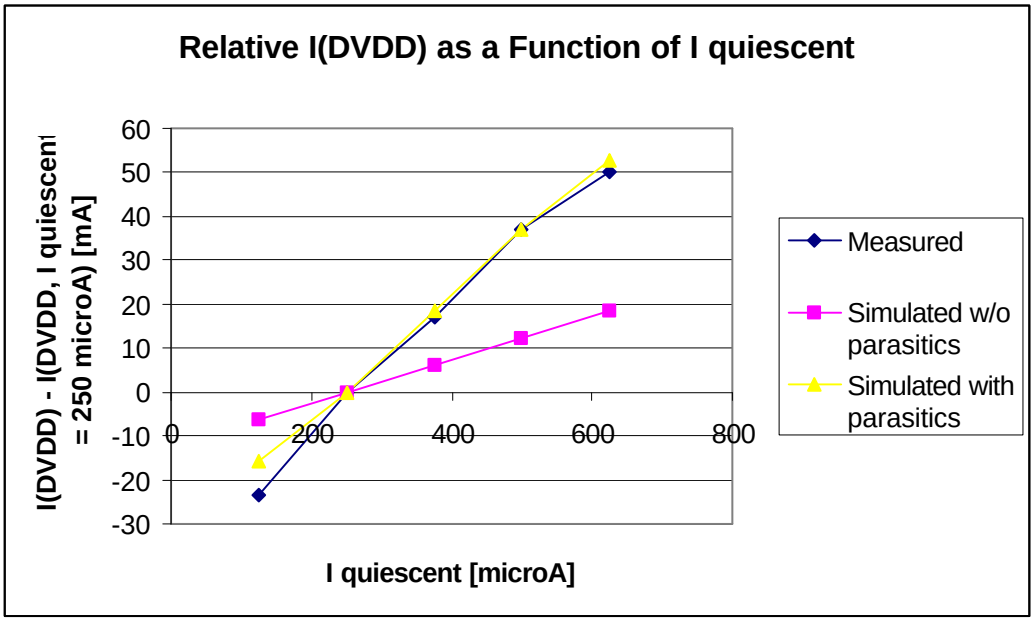

Figure 37 Relative current consumption of DVDD as a function of the I quiescent current.

The simulation and the measured relative currents agree very well.

17 Irradiation of the SVX4

Irradiation studies of the SVX4 chip were done at University of California Davis and are summarized below.

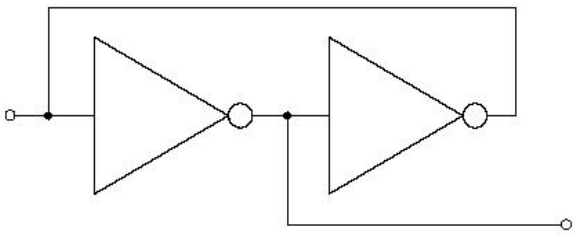

Figure 38 Cartoon picture of a CMOS memory cell.

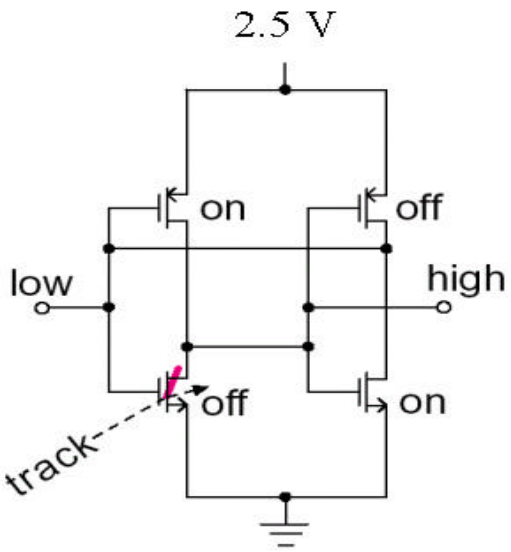

Figure 39 Cartoon schematic showing how a highly ionizing particle can flip a CMOS transistor from the off state to an on state.

17.1 Single Event Upsets

We were able to study Single Event Upsets, also known as SEUs, by setting the parameters of the chip and then monitoring the performance.

2 SEU test at UC Davis cyclotron
=== ==== == == ===== =========

Marc and Dave reported on the SEU tests. Overall, the tests went fairly smoothly. Three chips were irradiated with 63 MeV protons at fluences of up to 1.4 x 10^14 and total doses of maximum 19 Mrad. The main results are:

a The shift register is rather radiation hard.

Only 4 SEUs where detected. The SEU cross section is ~6.4x10^-17 cm^2, and to 95% confidence level it is in the range of 2.2x10^-17 to 1.6x10-^16 cm^2.

This is significantly harder than the SVX3d.
(The SVX3d SEU cross section obtained in the same beam varied from 2.4x10^-16 to 3.9x10^-15 cm^2 depending on the proton incident angle and shift register cell type.)

b) The shadow register is also OK.

Not a single SEU was detected in the shadow register (the chip ID and counter maximum bits were tested).

c) A first check of the analog performance shows that there is no major change in pedestal value, gain or noise. A few more studies will be done to check the fine details.

d) At ~15 Mrad a readout problem develops which was seen in two chips. (The third chip was only irradiated to ~10Mrad.)
The symptom is that events are written out which have the wrong chip and pipeline cell ID and/or are missing channel number 127.

One of the two chips was irradiated up to 19 Mrad. Then the problem got worse gradually. More channels were missing, others may appear twice or more times in the event record, the valid range of ramp pedestal gets reduced ...

Not all events are bad. The fraction of bad events also seems to increase with dose.

This readout effect will in particular be studied at the upcoming Co-60 irradiation at Sacramento at September 26/27.

Brad proposed to lower DVDD to check if that may reproduce the effect.

17.2 Cobalt 60 Testing

Marc briefly presented the measurement program for the Co-60 gamma ray irradiation at DMEA.

A single chip and a 4-chip CDF hybrid will be irradiated. Total doses of 15 Mrad are the goal. The dose rate in 2 Mrad/hour.

Chips will be operated and readout using the LBNL Patt DAQ system. Every hour a series of data sets, containing ~1000 events (all channels), will be written to disk. Data sets will vary in the value of VCAL, bandwidth, applied DVDD and AVDD, ADC ramp direction. From the data sample noise, pedestal structures, gain etc. will be calculated.

The frontend voltages Ncas and Bias will be monitored, as

well as the preamp input voltage level during reset.

The readout will be checked for the irregularities that where observed during the cyclotron proton beam irradiation.

After irradiation, and also a few days later, the chips will be tested again in order to learn about possible annealing effects.

A first compilation of the Co-60 irradiation data was shown. A single chip sitting on a PCB and a 4-chip CDF hybrid were exposed to the irradiation. They received doses of 17.9 Mrad and 15.8 Mrad respectively.
Small pedestal shifts and corresponding shifts of the response of channels with charge injection were observed. No obvious change of noise was seen.

There is a wealth of data which we have to analyze. In particular it will be checked if the pedestal chip is due to a change in gain. There are runs with different bandwith and different times and values of charge injection and runs at various AVDD and DVDD values...

18 Deadtimeless Operation

Only CDF will be using the deadtimeless operation of the chip. There is a jumper on the SVX4 chip carrier that must be set in order to choose which mode the chip will operate in.

First dead-timeless studies started at LBNL. Chips on the standard chip carrier board and on a CDF hybrids were studied.

The general measurement principle is this: first the chip is in acquire mode and a L1A is sent at a fixed time at the begin of the acquire sequence (which happens to correspond to pipe line cell 46). Then digitize and readout mode follow. The first L1A is only needed in order to be able to put the chip in the three different data acquisition modes. The data thus read out are not of any interest here and are ignored.

In order to check possible pedestal chips related to transitions between the modes or due to control signals within the modes a second L1A is sent. The second L1A is moved bucket by bucket through the above chip modes. This is the data we are interested in, thus another digitize and readout sequence follows in order to plot the data as a function of L1A position (or time or FE clocks if you like).

This sequence of chip states (and L1A signals) is repeated a couple of times. For each bucket, the data marked by the second L1A signal is averaged

over....

The good news is that there is only one big structure which happens at the end of RO and moves all channels into saturation. We will find out what this really is.
On the other hand there are a number of small structures. So far we identified a pedestal shift when the comparators fire (in the digitize sequence), and another one at the end of readout due to the simultaneous comp_rst, ramp_rst and rref_sel signals... . Also several pipeline errors occur after the big spike at the end of RO.

19 Specifications of the SVX4

We went through and measure all the specifications of the SVX4 chip and they are listed below. Various data was collected by the designers of the chip and that data is listed in tabular format.

	Design Value	Measured Value
Gain	3 mV/fC	4.4 mV/fC
Gain Uniformity	5% or better	Confirmed
External Load Capacitance	10 pF to 50 pF	Confirmed
Risetime 0-90%	Adjustable in 60-100ns for any load	Confirmed
Risetime Adjustment	4 bits	Confirmed
Noise (ENC)	<2000 e for 40 pF load	2025 e for 69 ns
Linearity	Linear response for pulses up to 20 fC	Confirmed
Dynamic Range	>200 fC	Confirmed
Reset + Settling Time	<1μs for any initial condition	100 ns + 1 FE clk
Calibration Injection	40 fF internal cap	25fF

Table 9 Preamplifier specifications and the measured values.

	Design Value	Measured Value
Voltage Gain	3-5	3.4
Gain Uniformity	5% channel to channel	Confirmed
Risetime 0-90%	10 ns to 40 ns	Confirmed
Noise (ENC)	<500 e	Confirmed
Linearity	Linear response for pulses up to 20 fC	Confirmed
Dynamic Range	>40 fC at preamp input	Confirmed
Reset Time	<20 ns for any allowed condition	25 ns, okay

Pedestal Uniformity	<500 e at preamp input (channel to channel) <1000 e at preamp input (cell to cell)	Confirmed Failed!

Table 10 Pipeline specifications and the measured values.

	Design Value	Measured Value
Ramp Rate Trim Bits	3 bits	Confirmed
Ramp Linearity	0.25% for rates between 0.1 and 1 V/μs	Confirmed
Counter	8 bit Gray code, 106 MHZ rate	Confirmed
Differential Non-Linearity	<0.5 LSB	Confirmed

Table 11 ADC specifications and measured values.

	Design Value	Measured Value
Current Source Range	2.5 mA to 17.5 mA in 2.5 mA steps	Confirmed
Rise and Fall Times	>2ns and < 4ns	Confirmed
Bi-directional	All bus pads are I/O for D0	Confirmed
Single Ended Use		Confirmed

Table 12 Data Output Driver specifications and measured values.

	Design Value	Measured Value
BN/TN Modes	Only Active in Digitize Mode	Confirmed
Priority In/Out Modes	Configuration register input/output Priority passing during Readout Cycle Priority Out High during	Confirmed Confirmed Confirmed

	Digitize Cycle	

Table 13 Top Neighbor and Bottom Neighbor (TN/BN) specifications and measured values.

	Design Value	Measured Value
Calibration Injection Mask	128 bits	Confirmed
Preamp Bandwidth	3 bits	Confirmed
Ramp Rate	3 bits	Confirmed
Preamp Current	2 bits	Confirmed
Pipeline Write Current	2 bits	Confirmed
Pipeline Read Current	2 bits	Confirmed
Pipeline Depth	6 bits	Confirmed
Driver Current	2 bits	Confirmed
Threshold for Sparsification	8 bits	Confirmed
Counter Modulo	8 bits	Confirmed
Chip ID	7 bits	Confirmed
Dynamical Pedestal Subtraction	1 bit	Confirmed
Read Neighbors	1 bit	Confirmed
Read Channel 63	1 bit	Confirmed
Read Channel 127	1 bit	Confirmed
Pedestal Adjustment	3 bits	Confirmed
Reversed Polarity	1 bit	Confirmed

Table 14 Configuration Register with all functionality.

20 Proposed changes to the SVX4

The have been several proposed changes to the SVX4 that will be implemented in the next version of the chip. We discuss each one of them below.

20.1 Pullup or Pulldown on D0 mode pad

As stated in the introduction, the chip can be operated in two modes – D0 and CDF. Presently, the pad that selects this mode, the D0 mode pad, is floating. That means the

pad must be wire bonded to DVDD for the chip to be operated in D0 mode or it must be wire bonded to GND for the chip to be operated in CDF mode. It must be decided which mode will be the default operating mode of the chip which is done by placing either a pullup resistor to DVDD or a pulldown resistor to GND inside the chip.
Suppose that the wire bond connecting this pad to the hybrid fails, the chip will go into the default mode. If there was a pullup resistor, the chip will go into D0 mode. If there was a pulldown resistor the chip will go into CDF mode. It must be decided between the two collaborations which configuration this pad will default to. If the default is chosen to be D0 mode, then no wire bond will be needed for this pad. If the default is chosen to be CDF mode, then this pad will have to be wire bonded to DVDD to operate in D0 mode.

20.2 Pullup on USESEU pad

As stated in the introduction, the current version of the chip allows the user to decide which register to initialize the chip from. If the USESEU pad is wire bonded to ground, the operating parameters will be taken from the shift register. If the USESEU pad is wire bonded to DVDD, the operating parameters will be taken from the SEU register (or shadow register). Since the SEU register design has been confirmed to work and is radiation hard, this will be the default configuration. As a side note, the shift register cannot be used due to the Catch-22 effect of the driver currents changing as you download the initialization bit stream.

20.3 Pullup of Bit 7 on Chip ID

The chip ID only goes up to 128, so that leaves one extra bit in the 8 bit readout bus that is arbitrary. Since the SVX4 will be operating in sparse mode in the experiment, it becomes possible to misinterpret the chip ID as a channel number followed by its data because the length of the readout is unknown. By pulling up bit 7 of the chip ID output, the channel IDs are effectively changed to 129-255. Therefore it becomes impossible to misinterpret the chip ID with a channel number. Note: there are 128 channels that can have data values from 0-255. The channel numbers are hardwired in the readout.

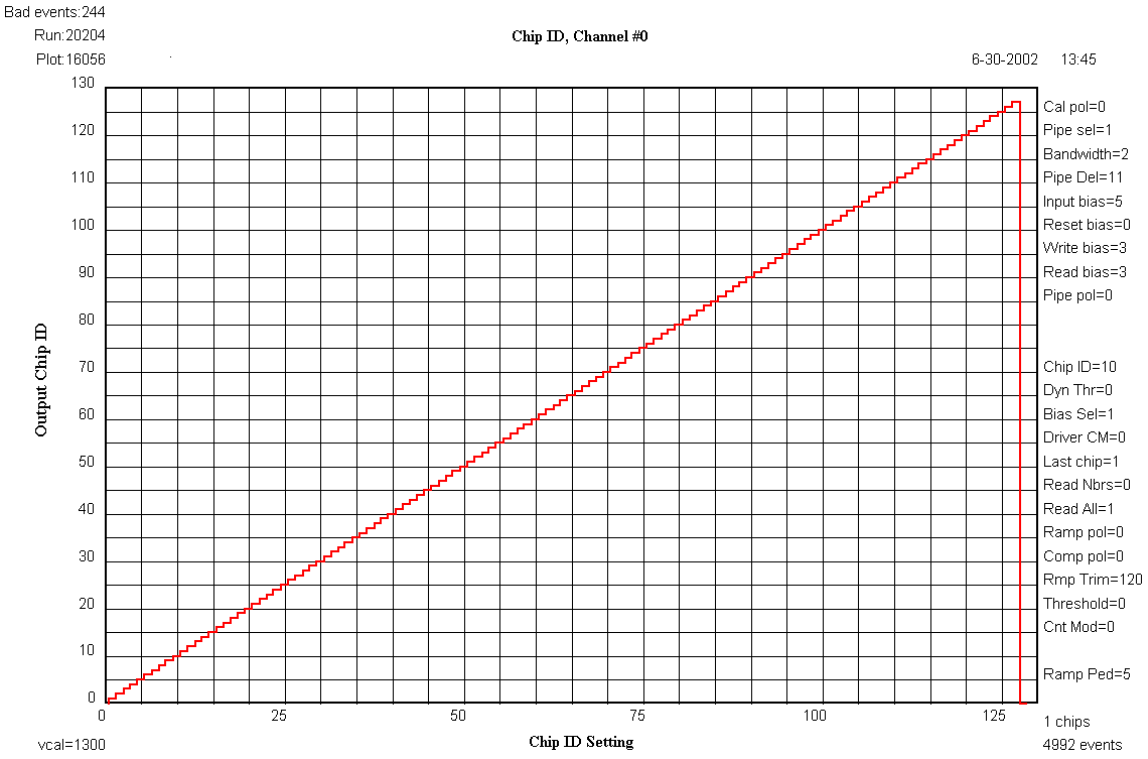

Figure 40 Histogram showing a chip ID scan. The x-axis shows the value of the chip ID that was downloaded to the chip and the y-axis shows that value of the chip ID that was read from the chip in readout cycle.

20.4 Pull down of Bit 6 and Bit 7 on Cell ID

Since there are only 47 pipeline cells that can be used, the status word does not need all 8 bits similar to the chip ID. It is easier to read out the pipeline cell number directly into the DAQ system if the upper two non-used bits are pulldown or always zero.

20.5 Hardwiring PRIOUT driver strength bits

During early testing of the SVX4, it was noticed that when you tried to initialize the SVX4 chip using the shift register as opposed to the USESEU register, you were not able to download the parameters to the chip. This is because if you choose to use the shift register to initialize your chip, while you are downloading you are effectively changing the bits of the PRIOUT driver which will skew the downloading sequence.

20.6 Adding 2 more bits to the shift register

The current length of the download of the SVX4 is 189 bits. This does not fall with a word boundary, or said another way, 189 is not divisible by two. Adding an additional two bits makes the downloading length fit within a word boundary. These additional bits can also be used for additional functionality if needed. One such use is described below.

20.7 Changing latching scheme for ADC control in D0 mode

There is a design flaw in the SVX4 that was not properly observed in the simulation. The design flaw is that the front end clock is not gated during the digitization in D0 mode. This makes it impossible in D0mode to place the pipeline cell onto the write amplifier. There is a work around that is complicated which will effect the rate and deadtime of the chip when operating in D0mode. This modification will make the digitization process easier and will decrease the amount of deadtime of the chip.

20.8 Adding a V cal switch

Because the D0 hybrids or purple card cannot source enough current to set the V cal voltage properly, it has become attractive to hae a switch that will turn of the V cal voltage and allow an external pulser to be used to calibrate the system. The extra 2 bits that are added as discussed above, one bit will have the functionality to turn of V cal.

20.9 Adding on-chip decoupling capacitors to BIAS

The use of on chip bypassing for the power AVDD has led to the proposal of the idea to bypass the BIAS by on chip bypassing as well. There is enough space on the chip in order to do this. Essential you are construction transistors below the pipeline which total up to a microfarad in capacitance and function as you power bypass capacitor. This removes one extra component from the hybrid that you will need.

20.10 Layout and Design changes to the ADC comparator and improvements.

The have been many suggesting to improving the design of the ADC. The biasing scheme of the comparators will be modified in order to improve the response time of the comparators. Currently, the pedestal can only be lowered to 60 ADC counts due to the analog delay in the circuit. The analog delay in the circuit is there to produce a baseline in DPS mode in order to do noise studies in that mode.

20.11 Modifying the LVDS receiver

It has been found that lowering or raising the I quiescent current will prevent the receiver from working properly. This is because there is a size mismatch in the p and n fets that make up the receiver. By increasing the size of the p fets of the receivers will correct this problem and the chip will work properly for a large value of I quiescent. There is also a discussion to remove the receiver from I quiescent and bias it by itself through another transitor.

20.12 Increasing the width of the Pipeline Metal

As discussed above, there is a slope in the pipeline which has no solid explanation. It is believed to be coming from the bus in the pipeline which has become resistive and therefore it slightly alters the gain for each individual channel of the pipeline.

20.13 NMOS Guard Rings

It was found that the NMOS transistors do not have a PMOS guard ring which is standard practice for making electronic circuitry radiation hard.

21 Traps and Pitfalls

Usually, the acquire knowledge and experience from chip testing is passed down by word of mouth. This section is an attempt to document the experience gained while testing the prototype SVX4 chips. This section may lack coherence and fluidity, but that is a small price to pay in order to have a list of the nuances learned from early testing.

21.1 Pipeline Cell 63

It was found during testing that pipeline 63 was being returned from the chip while testing. There were a variety of reasons that this occurred. First, it the mode lines are not properly set to enter the acquire cycle, it appeared that the pipeline logic simply defaulted to pipeline 63. Second, this could also occur if no Level 1 accept was received by the chip while in the acquire cycle.

21.2 L1A in conjunction PRD2

It was found during testing by the designers that in deadtimeless mode, if the Level 1 accept was received while the PRD2 was being issued, the chip sometimes would not see the Level 1 accept. This is because there are strict time constraints that must be adhered to during this mode of operation. The PRD2 pulse must be wide enough so there is no conflict between the rising and falling edges of the pulses.

21.3 Incomplete readout or zeros in the data

It was also found that sometimes the data read back from the chip while in read all mode did not read back all the channels or that the data contained a majority of zeros. Sometimes the data read back was incomplete.
It was found that it was necessary to be completely sure about the number of clock edges that the chip received in digitization mode. Sending 256 edges is not enough to complete the digitization process. It is necessary to send two extra close pulsed in order to synchronize the output latch with the clock. If there two clock pulses were not sent it was possible that the data value from the channel would not be latched in.
Also, a timing violation in the readout FIFOs occurs in the chip if the ramp and comparator reset setting were not maintained at the proper digital level during the complete digitization cycle. It was found that either the readout was incomplete or readout would fail completely. Once the comparator reset and ramp reset are released, they must stayed in that condition until the readout cycle occurs.

22 The Future

It has been decided to take two approaches for the next submission of the SVX4. We are going to submit three different designs: original, minimal changes, and maximal changes. The original design will be submitted as a baseline to study the process variations. The minimal changes is a plan to make the minimal changes that will correct the pedestal

bowing with low risk. The maximal changes model incorporates redesigning of the ADC and other features.

23 Conclusions

We have extensively tested the SVX4 chip using the Stimulus test stand. We have confirmed the specifications that have been specified and have documented and given a possible explanation for the specifications that have not been satisfied.
The next chip submission will address the performance issues with a slight modification of the ADC.